Amplifying Randomized Encodings & Applications

Pouria Fallahpour[1], Alex B. Grilo[1], Garazi Muguruza[2]
, and Mahshid Riahinia[3*]

[1] Sorbonne Université, CNRS, LIP6, France
pouria.fallahpour,alex.bredariol-grilo@lip6.fr
[2] SURF, The Netherlands
garazi.muguruzalasa@surf.nl
[3] DIENS, École Normale Supérieure, CNRS, Paris, France
mahshid.riahinia@ens.fr

Abstract. A randomized encoding for a promise problem is a randomized reduction whose outcome distribution on input x can be simulated within some distance d, called privacy, using only one bit of information about x: whether it is a YES or NO instance. The encoding is one-sided if this property holds only for YES instances. Our main contribution is showing that (classical and quantum) one-sided randomized encodings have privacy and correctness amplification: any problem with a randomized encoding with privacy $1-1/\mathsf{poly}(n)$ and error $1/2-1/\mathsf{poly}(n)$, also has an encoding with negligible privacy and error. We then use this result to obtain the following statements:

- We show that NISZK, the class of problems that have a non-interactive (statistical) zero-knowledge proof, has strong zero-knowledge amplification, i.e., every problem with a non-interactive zero-knowledge proof with zero-knowledge error $1-1/\mathsf{poly}(n)$, also has a non-interactive zero-knowledge proof with negligible zero-knowledge error. This solves a problem open since the work of Goldreich, Sahai, and Vadhan (CRYPTO '99).
- The worst-case hardness of a problem that has a perfect (i.e., zero error) one-sided encoding implies one-way functions (OWFs). And if the encoding is quantum, it implies one-way state generators (OWSGs). We then conclude that by removing the error from one-sided randomized encodings, the worst-case hardness of the complexity class SZK becomes sufficient for the existence of OWFs.
- *Weak* and *imperfect* indistinguishability obfuscation (iO) implies OWFs assuming that the Polynomial Hierarchy does not collapse to its third level. Here, weakness means that the computational distance between the obfuscated circuit and the original circuit is $1-1/\mathsf{poly}(n)$, and imperfectness means that the error is $1/2-1/\mathsf{poly}(n)$.

To achieve this privacy and correctness amplification, we study randomized encodings through the lens of lossy reductions (Ball et al. ITCS 2020). In particular, we introduce an extended and flexible notion of lossy reductions, and show that it is equivalent to randomized encodings. This equivalence serves as the base of our main results and might be of independent interest.

1 Introduction

For a function f and an input x, can one encode x into a string z_x that reveals nothing about x beyond $f(x)$? Of course, the value $f(x)$ is itself a valid candidate for the encoding, so the question becomes non-trivial if we require that computing the encoding be simpler than computing f. It can be shown that deterministic encodings cannot generally satisfy this requirement [5]. However, randomness changes the picture: for a wide range of complex functions f, there exist randomized encodings that are far simpler to evaluate than f itself.

*Part of this work was done when the author was visiting IRIF, Paris.

The gap between the complexity of $f(x)$ and that of $x \mapsto z_x$, together with the privacy of z_x, makes randomized encodings a rich tool for complexity theory as well as cryptography.

Ishai and Kushilevitz [39] showed that every function computable by a polynomial-size non-deterministic branching program ($\mathsf{NL/poly}$) has a randomized encoding that can be evaluated as a polynomial of degree 3. Consequently, securely evaluating any such function reduces to securely evaluating degree-3 polynomials, which yields round-efficient secure-computation protocols for all $\mathsf{NL/poly}$. Applebaum, Ishai, and Kushilevitz [6] pushed this further, and replaced the degree-3 polynomials by constant-depth circuits of locality 4 (denoted by NC^0_4). The evaluation of any function in $\mathsf{NL/poly}$, including candidate pseudorandom generators (PRGs) and pseudorandom functions (PRFs), can thus be brought all the way down to NC^0_4, which provides an evidence to the utility of randomized encodings. Brakerski and Yuen [18] introduced a quantum analogue, i.e., quantum encodings, and showed applications in quantum zero-knowledge proofs.

More generally, for any promise problem Π one can study randomized encodings of its characteristic function χ_Π, which on input x outputs 1 if $x \in \Pi_Y$ and 0 if $x \in \Pi_N$. In this terms, the discussion above implies that the class of problems with NC^0_4 encodings contains $\mathsf{NL/poly}$. What if we ask only that the encoding be computable in polynomial time? The resulting class has already been characterized [7]: $\mathsf{RE}[d(n), \mu(n)]$ denotes the class of all problems Π for which there exists a polynomial-time algorithm Enc (the randomized encoding) that satisfies the following two properties:

- (d-Privacy) There exists a polynomial-time algorithm Sim such that for all n-bit valid instances x:
$$\|\mathsf{Sim}(1^n, \chi_\Pi(x)) - \mathsf{Enc}(x)\| \leq d(n) \ ,$$
where $\|X - Y\|$ denotes the statistical distance between two distributions X and Y.

- (μ-Correctness) There exists a (possibly unbounded) algorithm Dec such that for all n-bit valid instances x:
$$\Pr\left[\mathsf{Dec}(\mathsf{Enc}(x)) \neq \chi_\Pi(x)\right] \leq \mu(n) \ .$$

If the existence of the simulator is only guaranteed for the YES instances, the encoding is called *one-sided* and the corresponding class of promise problems is denoted by $1\mathsf{RE}[d(n), \mu(n)]$. When $d(n), \mu(n) = \mathsf{negl}(n)$, we drop these arguments for simplicity.

The class RE trivially contains BPP, that is the class of problems admitting probabilistic polynomial-time solvers with bounded error. However, no explicit non-trivial encoding is known for any problem beyond BPP. There is evidence that the containment is nonetheless strict: the work of [1] showed that under suitable computational hardness assumptions, RE must be strictly larger than BPP. On the flip side, Applebaum and Raykov [7] showed that $1\mathsf{RE}$ is equivalent to NISZK, the class of all problems admitting a non-interactive (statistical) zero-knowledge proof system. Since $\mathsf{RE} \subseteq 1\mathsf{RE}$, this bounds both classes. Despite these results, RE and $1\mathsf{RE}$ remain little understood, and yet important objects of study. Randomized encodings in constant depth gave us cryptography of remarkably low complexity; in the polynomial-time regime, separating RE from BPP would give us cryptography of another kind, namely OWFs [7]. Moreover, the close connection between polynomial-time randomized encodings and non-interactive zero-knowledge proofs makes them a natural tool for studying proof systems. In this work, we study the properties of randomized encodings and show various applications to complexity theory and cryptography.

1.1 Our Contribution

In this work, we introduce a new technique for *privacy and correctness amplification* of $1\mathsf{RE}$. More precisely, we prove the following lemma:

Key Lemma (Privacy & Correctness Amplification of 1RE)**.** *For any sufficiently large $n \in \mathbb{N}$, it holds that*

$$1\mathsf{RE}\left[1 - 1/\mathsf{poly}(n), 1/2 - 1/\mathsf{poly}(n)\right] \subseteq 1\mathsf{RE}\left[\mathsf{negl}(n), \mathsf{negl}(n)\right] ,$$

up to some polynomial-size non-uniform advice. Moreover, if the error of the left-hand side is zero, the error of the right-hand side will also be zero.

The above result is obtained by studying the *lossy* properties of randomized encodings. More precisely, we show that the class 1RE can be equivalently characterized by a quantitative measure that we call *sparse lossiness*. Let us first define this measure formally. Our starting point is the notion of *lossy reductions*, introduced by Ball et al. [9]. In their definition, a reduction R for a promise problem Π is $\lambda(n)$-lossy if $I(X; R(X)) \leq \lambda(n)$ for *all* distributions X supported on $\{0,1\}^n$. We relax this requirement as follows. Call X a *sparse uniform* distribution if it is uniform over a set of $\mathsf{poly}(n)$ strings that lies entirely in $\Pi_Y \cap \{0,1\}^n$ or entirely in $\Pi_N \cap \{0,1\}^n$. We then say that a reduction R is $\lambda(n)$-*sparsely lossy* if $I(X; R(X)) \leq \lambda(n)$ for all such X. In other words, we only require R to lose information on sparse uniform distributions of YES or NO instances. We let $\mathsf{SL}[\lambda(n), \mu(n)]$ denote the class of all problems that have $\lambda(n)$-sparsely lossy reductions with correctness error $\mu(n)$. Moreover, the reduction is said to be *one-sided sparsely lossy* if it is sparsely lossy only on sparse uniform distributions over YES instances; we denote the corresponding class of problems by $1\mathsf{SL}[\lambda(n), \mu(n)]$. This refined definition allows us to prove the following theorem:

Theorem 1. *For any error function $\mu(n) \leq 1/2 - 1/\mathsf{poly}(n)$ the following statements hold:*

1. *One-sided randomized encodings with inverse-polynomial privacy are equivalent to one-sided sparsely lossy reductions:*
   - *$\forall d(n) \leq 1 - 1/\mathsf{poly}(n)$, we have $1\mathsf{RE}\left[d(n), \mu(n)\right] \subseteq 1\mathsf{SL}[\mathsf{negl}(n), \mathsf{negl}(n)]$. Moreover, if $\mu(n) = 0$, the error of the right-hand side is 0.*
   - *We have $1\mathsf{SL}[O(\log n), \mu(n)] \subseteq 1\mathsf{RE}\left[\mathsf{negl}(n), \mu(n)\right]$.*

2. *Randomized encodings are partially equivalent to sparsely lossy reductions:*
   - *We have $\mathsf{RE}[\mathsf{negl}(n), \mu(n)] \subseteq \mathsf{SL}[\mathsf{negl}(n), \mu(n)]$.*
   - *We have $\mathsf{SL}[O(\log n), \mu(n)] \subseteq \mathsf{RE}\left[1 - 1/\mathsf{poly}(n), \mu(n)\right]$.*[1]

The statements hold up to some polynomial-size non-uniform advice.

One can verify that the first item of this theorem implies our Key Lemma. The conceptual connection between randomized encodings and lossy reductions was already observed by [9]. In Theorem 1, we make this connection precise and give an exact quantitative relation between the two.

Quantum Setting. The techniques we use to prove Theorem 1 are generic, and largely independent of whether the underlying objects are quantum or classical. This allows us to obtain a quantum version of the statement. More precisely, let QRE and 1QRE be the quantum analogues of RE and 1RE, where the encoding and the simulator are allowed to be quantum algorithms. The inputs, which are instances of computational problems, remain classical, and statistical distance is replaced by trace distance. With these definitions, Theorem 1 and Key Lemma hold verbatim for QRE and 1QRE.

[1] The inclusion can be extended to sparsely lossy f-reductions, the reductions that take multiple instance as input such as an Or reduction. More details can be found in Section 3.

Remark 1 (Compressing reductions and f-distillations). Compressing reductions have been studied in parameterized complexity. It is mostly useful to consider compression of multiple instances of a problem, therefore, f-reductions appear frequently in this context. A Π is said to have strong f-compression if there exists a polynomial-time f-reduction from Π to itself with negligible error such that it maps an m-tuple of n-bit instances $(x_1, \cdots, x_m)$ to an instance y that has bit size $O(m \log(n))$. Such a reduction is also known as f-distillation. Fortnow and Santhanam [30] show that NP does not have strong OR-compression unless Polynomial Hierarchy collapses. Drucker extended this result to strong AND-compressions. Any Π that has a strong f-compression (or f-distillation) belongs to $\mathsf{SL}_f[O(\log n)]$ and by our results to $\mathsf{NP/poly} \cap \mathsf{coNP/poly}$. In other words, NP does not have any f-compression unless Polynomial Hierarchy collapses. This can be extended to weaker forms of f-compressions. More precisely, using our techniques, one can show that if refuting 3SAT is exponentially hard, then does not have any f-compression that maps mn bits to $mn^{1-\varepsilon}$ bits for some constant $\varepsilon < 1$. These two extensions might be of independent interest in parameterized complexity and kernelization theory. These results are known as distillation lemmas that are useful for showing that some problems does not have polynomial-size kernels. Our result implies that $IC_f[O(\log n), \mathsf{negl}(n)] \subseteq \mathsf{RE}[1 - 1/\mathsf{poly}(n), \mathsf{negl}(n)] \subseteq \mathsf{AM}$ and therefore $\mathsf{NP} \notin \mathsf{IC}_f[O(\log n), \mathsf{negl}(n)]$ for every permutation-invariant f, unless Polynomial Hierarchy collapses. This might be of independent interest in parameterized complexity and kernelization theory.

Despite its simple statement, the Key Lemma is a powerful tool for obtaining further results. In particular, we use it to obtain the following applications.

Application 1: Zero-knowledge Amplification of NISZK. Recall that SZK is the class of problems who have a zero-knowledge interactive proof system. More precisely, in a zero-knowledge proof system for Π, an unbounded prover interacts with a polynomially-bounded verifier in order to prove that a given input $x \in \Pi_Y$ is a YES instance, without revealing any extra information about x beyond this membership. Three quantities $(c(n), s(n), \varepsilon(n))$ measure the quality of the proof system:

- **Completeness**: For any $x \in \Pi_Y \cap \{0,1\}^n$, the verifier accepts with probability at least $1 - c(n)$.
- **Soundness**: For any $x \in \Pi_N \cap \{0,1\}^n$ and every possibly malicious prover, the verifier accepts with probability at most $s(n)$.
- **Zero-Knowledge**: There exists a polynomial-time algorithm Sim such that for every $x \in \Pi_Y \cap \{0,1\}^n$, the distribution of $\mathsf{Sim}(x)$ is within $\varepsilon(n)$ statistical distance from the honest transcript of the protocol.

In its non-interactive variant, denoted NISZK, the interaction between the prover and the verifier is limited to a single message, and both parties have access to a shared randomness that is sampled from an *a priori* fixed distribution, known as the common reference string.

In the usual definition of SZK (respectively NISZK), all three parameters are set to be negligible. Moreover, as long as $\varepsilon(n)$ is negligible, any protocol with $c(n)+s(n) \leq 1-1/\mathsf{poly}(n)$ can be amplified into one where all three parameters are negligible, and in fact exponentially small [34, 56]. Amplifying the zero-knowledge property itself, meaning reducing $\varepsilon(n)$ when it is not already negligible, appears to be more challenging. The one known result in this direction is due to Goldreich, Sahai, and Vadhan [34], who considered a *weak* zero-knowledge property for NISZK. This property requires that for every inverse polynomial $1/p(n)$, there exists a polynomial-time simulator whose output is within statistical distance $1/p(n)$ of the honest transcript. They showed that any such protocol can be amplified to one with a single

simulator of negligible error. A similar result holds for SZK [56]. Yet, the following question has remained unanswered ever since:

Question 1. Can the zero-knowledge guarantee of SZK *or* NISZK *protocols be amplified from inverse-polynomial to negligible error?*

We prove the following theorem which answers this question in the affirmative for NISZK:

Theorem 2. *For every* $c(n), s(n), \varepsilon(n)$ *satisfying* $2s(n)+c(n)+\varepsilon(n) = 1-1/\mathsf{poly}(n)$*, it holds that* $\mathsf{NISZK}[c(n), s(n), \varepsilon(n)] \subseteq \mathsf{NISZK}$ *(up to some non-uniform polynomial-size advice).*

In particular, when $c(n) = s(n) = \mathsf{negl}(n)$ and $\varepsilon(n) = 1 - 1/\mathsf{poly}(n)$, the theorem above yields a zero-knowledge amplification for NISZK from inverse-polynomial to negligible.

Application 2: OWFs and the Worst-Case Hardness of SZK. One-way functions (OWFs) are fundamental cryptographic tools whose existence can be viewed as the minimal assumption required for cryptography. It is not known whether OWFs unconditionally exist or not. Instead, their existence is often based on the worst-case hardness of concrete computational problems such as DLOG [27], RSA [55], lattice and code-based problems [2,4,48,54] and more. While a polynomial-time algorithm for any of these problems would have considerable implications for cryptography, the widespread belief of their worst-case hardness is based less on *structural* or *complexity-theoretic* evidence and more on the failed *algorithmic* attempts in directly solving them. In the broader sense, relying on just a few specific problems remains limiting. A more robust approach is to look for OWFs based on the worst-case hardness of a large class of problems. These types of connections provide stronger evidence for the existence of OWFs.

For OWFs to exist, it is necessary that $\mathsf{P} \neq \mathsf{NP}$. Whether this is sufficient is arguably one of the most important questions in the theory of cryptography. In a series of works [3,14,15,29], it has been shown that there are black-box barriers to build OWFs from worst-case hardness of NP. In fact, in the non-uniform setting where the algorithms have access to advice, some of these barriers also hold even for SZK-complete problems. Recall that SZK is the class of problems that have statistical zero-knowledge interactive proofs and is inside $\mathsf{NP} \cap \mathsf{coNP}$ (in the non-uniform setting). On the other hand, many of the widely used assumptions such as lattice-based or group-based ones [2, 27, 48, 54, 55] belong to SZK.[1] This suggests that SZK may provide sufficient hardness for building OWFs. Ostrovsky [52] showed that the the average-case hardness of SZK yields OWFs. Later, Ostrovsky and Wigderson [53] could show that the worst-case hardness of SZK implies the (seemingly weaker) auxiliary-input OWFs.[2] A positive result relating OWFs to the worst-case hardness of SZK would therefore sheds more light on the sufficient hardness for OWFs.

Question 2. Can we show that the worst-case hardness of SZK *suffices for the existence of OWFs?*

We show that an error removal for 1RE, even with loose privacy, is sufficient to prove that the worst-case hardness of SZK implies the existence of OWFs. More formally,

Theorem 3. *If* $\mathsf{1RE}[\mathsf{negl}(n), \mathsf{negl}(n)] \subseteq \mathsf{1RE}[1 - 1/\mathsf{poly}(n), 0]$ *(up to some non-uniform polynomial-size advice), and* SZK *is worst-case hard, then OWFs exist.*[3]

[1] Yet, none of them are complete for SZK.

[2] This is different from OWFs with non-uniform advice. We refer the reader to the cited paper for more details.

[3] We note that by the Key Lemma, the privacy can always shrink to negligible, but we keep $1 - 1/\mathsf{poly}(n)$ privacy in the statement for expressiveness.

We expect that an error removal algorithm for $\mathsf{1RE}$ exists. Without considering the privacy, a polynomial-time randomized reduction can be derandomized via the textbook repeat-then-majority procedures. Here, the challenge is to manipulate the randomness in order to not completely lose the structure (privacy).

Application 3: OWFs from Weak and Imperfect iO. An Indistinguishability Obfuscation (iO) is an algorithm that takes as input a circuit $C : \{0,1\}^n \to \{0,1\}^*$ and outputs an obfuscated circuit such that (I) the latter has the same input-output behavior as the original circuit except with probability $\varepsilon = \mathsf{negl}(n)$, and (II) for every two functionally equivalent circuits C, C', their obfuscated variants are computationally distinguishable with probability at most $\tau = \mathsf{negl}(n)$. Numerous works are dedicated to study the existence of obfuscation, e.g., [10, 11, 31, 40, 41, 62], and its importance in cryptography as well as complexity theory, e.g., [8, 13, 19, 24, 25, 32, 35, 57]. In particular, Goldwasser and Rothblum [36] implicitly show that when $\tau \approx 1/10$ and $\varepsilon = \mathsf{negl}(n)$, iO implies OWFs unless the Polynomial Hierarchy collapses. Komargodski et al. [44] show that iO with $\tau = \mathsf{negl}(n)$ and $\varepsilon = 1/2 - 1/\mathsf{poly}(n)$ implies OWFs if $\mathsf{NP} \not\subseteq \mathsf{BPP}$. A simple argument aligned with that of [49] also yields OWFs from iO when $\tau = 1 - 1/\mathsf{poly}(n)$ and $\varepsilon = 0$, if $\mathsf{NP} \not\subseteq \mathsf{BPP}$. All these regimes relax one of τ or ε while keeping the other close to optimal. This raises the following question:

Question 3. Does iO imply OWFs when both τ and ε are simultaneously far from optimal?

In this regard, we prove the following theorem:

Theorem 4. *Assuming the existence of iO with parameters $\tau(n) = 1 - 1/\mathsf{poly}(n)$ and $\varepsilon(n) = 1/2 - 1/\mathsf{poly}(n)$ and that the Polynomial Hierarchy does not collapse to its third level, OWFs exist.*

Moreover, assuming the existence of iO with subexponential runtime and parameters $\tau(n) = 1 - 1/\mathsf{subexp}(n)$ and $\varepsilon(n) = 1/2 - 1/\mathsf{subexp}(n)$, and that refuting 3SAT *is exponentially hard, OWFs exist.* [1]

The assumption that refuting 3SAT is exponentially hard is known as non-uniform non-deterministic exponential time hypothesis (nuNETH). It is introduced by [21], and more formally states that 3SAT does not have a non-uniform reduction to coNP that runs in time subexponential in n, where n is the number of variables in the 3SAT instance. It has been further studied by [21, 23].

On the existence of iO. In fact, we prove a stronger result concerning the existence of iO. We show that unless the PH collapses to its third level, any obfuscation algorithm with error $1/2 - 1/\mathsf{poly}(n)$ must have a *statistical* distance negligibly close to 1 on at least one family of functionally-equivalent pair of circuits. Similarly, assuming nuNETH, any obfuscation algorithm with error $1/2 - 1/\mathsf{subexp}(n)$ must have a statistical distance exponentially close to 1.

Application 4: One-Way State Generators. We also obtain a weaker variant of Theorem 3 for *inefficient-verifier one-way state generators* (IV-OWSG). Roughly speaking, an IV-OWSG consists of a polynomial-time quantum algorithm $\mathsf{StateGen}$ that maps $k \in \{0,1\}^n$ to a quantum state ρ_k and a possibly unbounded algorithm Ver that, given k and a state ρ, either accepts or rejects. For instance, one way to instantiate Ver is to check whether ρ is close to ρ_k in trace distance. Moreover, it must be hard for any polynomial-time adversary to find a k', given multiple copies of ρ_k, such that $\mathsf{Ver}(k', \rho_k)$ accepts. When ρ_k is a pure quantum state for every fixed k, we say that the IV-OWSG is pure. Almost immediately by recycling the techniques developed for Theorem 3, we show that if $1\mathsf{QRE}_{\mathsf{pure}}[1 - 1/\mathsf{poly}(n), 0]$ is

[1] Note that iO with subexponential runtime should not be confused with subexponentially-secure iO.

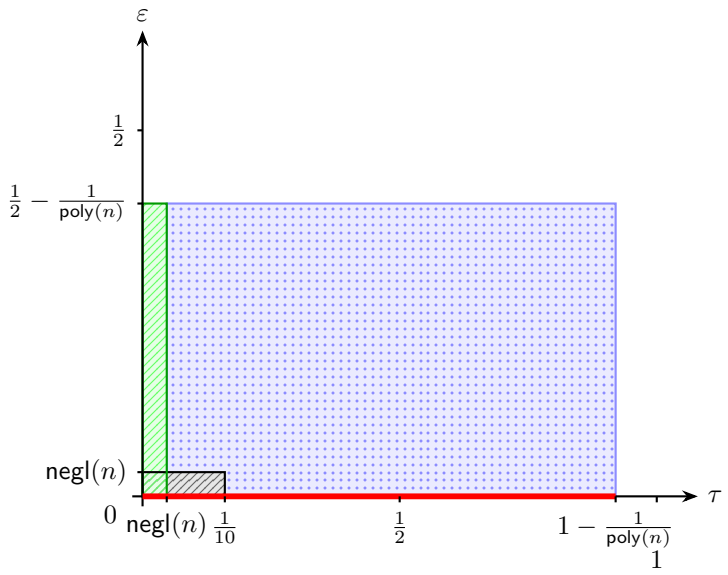

Fig. 1: The set of (ε, τ)-iO implying the existence of OWFs assuming that PH does not collapse. The grey and the green hatched regions are respectively the parameter sets of [36] and [44], and the bold red line corresponds to technique of [49]. The dotted blue area correspond to our result.

(quantum) worst-case hard, then pure IV-OWSGs exist. Here $1\mathsf{QRE}_{\mathsf{pure}}$ refers to those quantum encodings whose $\mathsf{Enc}(\cdot; r)$ function returns a pure state for any fixed possible classical randomness r.

1.2 Technical Overview

In this section, we briefly present the core technical tools that we use. We first discuss the proof of Theorem 1, the main technique to obtain the other results. We then sketch the proofs of Theorems 2, 3 and 4. Almost all our techniques immediately carry over to the quantum setting, leading to the quantum results. To avoid the risk of being redundant, we have not included the overview of these proofs in this section.

Proof sketch of Theorem 1 (Randomized encodings and sparsely lossy reductions). In what follows, we sketch the proof of Theorem 1 separately for the two directions of the inclusions.

From sparse lossiness to encodings. The difficulty in going from SL to RE lies in extracting a privacy simulator out of the sparse lossiness property. To make this happen, we resort to the *disguising lemma* by Drucker [28]. The proof works for general *f-reductions*, but for simplicity, we restrict the discussion to regular reductions. Let $\Pi \in \mathsf{SL}$ and R be a reduction from Π to some other problem that is sparsely lossy. The original variant of the disguising lemma states that if a mapping R is compressing, then for any subset $S \subseteq \{0,1\}^n$, there exists a sparse distribution D_S supported on S, such that for all $y \in S$, the distribution of $R(y)$ is statistically close to that of $R(D_S)$, where the closeness depends on the amount of compression achieved by R—therefore, if R is sufficiently compressing, then R *disguises* y. More precisely, $\mathbb{E}_{D_S}[\|R(y) - R(D_S)\|] \leq \delta^*$, where $\delta^* \approx 1 + 2^{-\lambda}$, if R compresses an n-bit instance to λ bits. Looking at it differently, since R is a reduction from a problem Π, we can set $S := \Pi_Y \cap \{0,1\}^n$ and conclude that $R(D_S)$ is close to $R(y)$ in average, for every YES instance y. In other words, $R(D_S)$ is a simulator for YES instances. Therefore, we obtain randomized encodings from compressing reductions. In other words, the class of problems with compressing reductions are a subset of those with randomized encodings.

While the original disguising lemma is stated for compressing reductions, the work of Ball et al. [9] extends it for lossy reductions. Our challenge is to extend it even further for sparsely lossy reductions.

In order to sketch our improvements of this lemma, we first briefly go over Drucker's proof. Consider the following simultaneous-move two-player game: Given a set S and $d \leq |S|$, Player 1 chooses a d-sparse distribution D supported on S, and Player 2 chooses an element $y \in S$. Define the payoff as $\|R(y) - R(D)\|$. Drucker first shows that if R is compressing, then for any strategy $\mathcal{Y}$ for choosing y, there exist a d-sparse distribution D, that is more precisely the uniform distribution over d samples of $\mathcal{Y}$, such that:

$$\mathbb{E}_{\mathcal{Y}^{\otimes d}}[\|R(y) - R(D)\|] \leq \delta^* \,. \tag{1}$$

By the minimax theorem, we can switch the quantifiers in above, implying that there exists a distribution $\mathcal{D}_S$ such that, for any choice of $y \in S$, it holds that $\mathbb{E}_{D\sim\mathcal{D}_S}[\|R(y) - R(D)\|] \leq \delta^*$. However, note that $\mathcal{D}_S$ is not guaranteed to be sparse. Drucker finally uses a result by Lipton and Young [47, Theorem 2] which roughly states that restricting the strategies of Player 1 to uniform strategies with support size $\ln(\#\{\text{choices of Player 2}\})/\gamma^2$ only changes the optimal expectation payoff with an additive factor γ. [1] This therefore *sparsifies* the support of $\mathcal{D}_S$. Conversely, it incurs a loss of at most an additive factor γ in the expectation bound in Equation (1) and obtains $\delta^* + \gamma$.

We relax the requirement on R and show that *sparse lossiness* of R suffices to obtain a similar result. More precisely, we show that if R loses information on input distributions that are uniform with support size roughly $1/\gamma^3$, then one loses nothing but an(other) additive factor γ in Equation (1). This relies on a double use of the result by Lipton and Young [47, Theorem 2]; we apply it once for Player 2 and once more for Player 1. Before using the minimax theorem, we restrict Player 2's strategies to be uniform distributions with support size roughly $1/\gamma^3$, and we choose $d \approx 1/\gamma$. By showing that Equation (1) remains correct even with this new restriction, we obtain an additive γ-approximation of the value of the game (first use of [47, Theorem 2] for Player 2). Following the minimax theorem, and sparsifying $\mathcal{D}_S$ (via a second use of [47, Theorem 2] for Player 1), we achieve the final upper bound $\delta^* + 2\gamma$. Note the constraint that y must have the same support as the distribution that it replaces. Our variant of disguising lemma with sparsely lossy reductions also extends to this setting (we refer to Appendix A for more details), and allows us to prove

$$\begin{aligned} \mathsf{SL}[O(\log n), \mu(n)] &\subseteq \mathsf{RE}[1 - 1/\mathsf{poly}(n), \mu(n)]/\mathsf{poly}(n)\,, \\ 1\mathsf{SL}[O(\log n), \mu(n)] &\subseteq 1\mathsf{RE}[1 - 1/\mathsf{poly}(n), \mu(n)]/\mathsf{poly}(n)\,. \end{aligned} \tag{2}$$

For the one-sided variant, we show how to reduce the privacy to negligible in the next parts.

From encodings to sparse lossiness. Assume that $\Pi \in \mathsf{RE}$. Recall that the privacy property of the randomized encoding puts an upper-bound on the statistical distance between the output distributions of the encoding $\mathsf{Enc}(x)$ and the simulator $\mathsf{Sim}(\chi_\Pi(x))$. Since Sim depends only on $\chi_\Pi(x)$ (i.e., whether x is a YES or NO instance), rather than the whole input x, intuitively, Enc should discard all information about its input beyond its membership in the language. Yet, the proof is not immediate. We sketch this proof below.

We first explain why Enc is sparsely lossy for the case where privacy d is negligible. We have to show that $I(X; \mathsf{Enc}(X))$ is bounded for every distribution X that is sparse, splitting and uniform, as required by the definition of the sparse lossiness, provided earlier in this section. For any two random variables X and Y, it holds that $I(X, Y) = \mathbb{E}_Y[D_{\mathrm{KL}}(p_{X|Y} \| p_X)]$ where D_{KL} is the KL-divergence. It is therefore sufficient to bound $D_{\mathrm{KL}}\left(p_{X|\mathsf{Enc}(X)} \parallel p_X\right)$. We do this by leveraging the *reverse Pinsker's inequality*; for any two distributions $p \ll q$ (i.e., p is absolutely continuous with respect to q), the reverse Pinsker's inequality bounds $D_{\mathrm{KL}}(p\|q) \leq$

[1] The same holds for Player 2.

$(\log e/\alpha_q)\cdot \Delta(p,q)^2$, where $\alpha_q = \min_x\{q(x) > 0\}$. Applying this bound and taking the expectation over $\mathsf{Enc}(X)$ yields an upper bound on the lossiness of Enc as follows: $I(X;\mathsf{Enc}(X)) \leq O(d^2/\alpha_X)$, where d denotes the privacy of the encoding. Since X is uniform over $\mathsf{poly}(n)$-many elements, it holds that $\alpha_X = 1/\mathsf{poly}(n)$. Therefore, if $d = \mathsf{negl}(n)$, we have that that $I(X;\mathsf{Enc}(X)) \leq \mathsf{negl}(n)$. In other words, Enc is a sparsely lossy reduction for Π as long as $d = \mathsf{negl}(n)$.

For extending this result to larger values of $d(n)$, we use the *XOR lemma* [56]. In a nutshell, XOR lemma takes as input a polynomial $q(n)$ and transforms a pair of circuits into a new pair whose statistical distance is that of the initial pair to the power of $q(n)$. Suppose that $d(n) = 1-1/\mathsf{poly}(n)$. Applying it to $(\mathsf{Enc},\mathsf{Sim})$ yields a new pair (P,T). A similar argument as in the previous paragraph gives us an upper bound on $I(X;P(X)) \leq O(d^{2q}/\alpha_X) = \mathsf{negl}(n)$, when $q(n)$ is a sufficiently large polynomial. Therefore, it follows that P is a sparsely lossy mapping for Π. However the correctness error of P is not necessarily negligible. In fact, one can show that $\Delta(P(x_Y),P(x_N)) \geq 1-2\mu$ for YES and NO instances x_Y and x_N, when $\Delta(\mathsf{Enc}(x_Y),\mathsf{Enc}(x_N)) \geq 1-2\mu$. In other words, the error remains the same (see Lemma 2.6 for details). In order to reduce the error, we consider the new amplified reduction $P^{\otimes\ell}$, where we repeat P for ℓ times. By the direct product lemma [56], for a sufficiently large polynomial $\ell(n)$, we get $\Delta(P^{\otimes\ell}(x_Y),P^{\otimes\ell}(x_N)) \geq 1-2\exp\bigl(-\ell(1-2\mu)^2/2\bigr) = 1-\mathsf{negl}(n)$ while the effect on the mutual information is $I(X;P^{\otimes\ell}(X)) \leq O(\ell\cdot d^{2q}/\alpha_X) = \mathsf{negl}(n)$. This gives us

$$1\mathsf{RE}[1-1/\mathsf{poly}(n),\mu(n)] \subseteq 1\mathsf{SL}[\mathsf{negl}(n),\mathsf{negl}(n)]\,. \tag{3}$$

For two-sided randomized encodings, i.e., when the simulator works on both YES and NO instances, the amplification no longer works. In the output of the XOR lemma, P (and similarly T) is not solely obtained from Enc and internally runs Sim as well. Let $\mathsf{Sim}(1)$ and $\mathsf{Sim}(0)$ be the simulators for YES and NO instances, respectively. The issue comes from the fact that we obtain different P circuits when we apply the XOR lemma to $(\mathsf{Enc},\mathsf{Sim}(1))$ and $(\mathsf{Enc},\mathsf{Sim}(0))$. This prevents from coming to agreement for the choice of the reduction. Therefore, this method only works for proving the sparse lossiness of one-sided encodings.

From negligible sparse lossiness to negligible privacy. When λ, the sparse lossiness quantity, is sufficiently small, we construct another encoding different from the one obtained by disguising lemma. Let $\Pi \in \mathsf{SL}[\lambda,\mu]$ for some $\lambda \ll 1$ and R be its corresponding sparsely lossy reduction. Moreover, take as advice $x_1^* \in \Pi_Y \cap \{0,1\}^n$ and $x_0^* \in \Pi_N \cap \{0,1\}^n$ (in the one-sided variant, only x_1^* suffices). The randomized encoding for Π is as follows:

$$\mathsf{Enc}(y) := R\,(y) \quad \text{and} \quad \forall b \in \{0,1\} : \mathsf{Sim}(b) := R\,(x_b^*)\ .$$

The correctness transfers immediately by definition. For computing the privacy, we proceed as follows. Fix an input y and let $b := \chi_\Pi(y)$. We must bound $\|\mathsf{Enc}(y) - \mathsf{Sim}(b)\|$ from above. Let X be a random variable that is y with probability $1/2$ and x_b^* otherwise. Note that X has a uniform distribution over 2 elements. Therefore, the sparse lossiness guarantee that $I(X;R(X)) \leq \lambda$. Let M be a random variable that is $\mathsf{Enc}(y)$ with probability $1/2$ and $\mathsf{Sim}(b)$ otherwise. Then, we have that

$$I(X;R(X)) = \mathbb{E}_X[D_{\mathrm{KL}}(p_{R(X)|X}\|p_{R(X)})] = \frac{1}{2}D_{\mathrm{KL}}(p_{\mathsf{Enc}(y)}\|p_M) + \frac{1}{2}D_{\mathrm{KL}}(p_{\mathsf{Sim}(b)}\|p_M)\,.$$

The right-hand side can be bounded from below by $\|\mathsf{Enc}(y) - \mathsf{Sim}(b)\|^2/(2\ln 2)$ according to Pinsker's inequality and the fact that $\|\mathsf{Enc}(y) - M\| = \|\mathsf{Sim}(b) - M\| = \|\mathsf{Enc}(y) - \mathsf{Sim}(b)\|/2$. Therefore, if $\lambda(n) = \mathsf{negl}(n)$, the privacy of the encoding would be negligible. In other words, we obtain

$$\begin{aligned} \mathsf{SL}[\mathsf{negl}(n),\mu(n)] &\subseteq \mathsf{RE}[\mathsf{negl}(n),\mu(n)]/\mathsf{poly}(n)\,, \\ 1\mathsf{SL}[\mathsf{negl}(n),\mu(n)] &\subseteq 1\mathsf{RE}[\mathsf{negl}(n),\mu(n)]/\mathsf{poly}(n)\,. \end{aligned} \tag{4}$$

Finally, combining Equations (2), (3), and (4) yield Theorem 1.

Proof of Theorem 2 (amplifying NISZK). The proof is a direct application of Key Lemma. Applebaum and Raykov implicitly (see [7, Theorme 7 & 8]) show that for every integer $k \geq 1$, it holds that $\mathsf{NISZK}[c, s, \varepsilon] \subseteq 1\mathsf{RE}[k(c+\varepsilon), (2s)^k + k(c+\varepsilon)]$. By taking $k = 1$ and imposing the condition $2s(n) + c(n) + \varepsilon(n) = 1 - 1/\mathsf{poly}(n)$ the former collapses to $1\mathsf{RE}[1 - 1/\mathsf{poly}(n), 1 - 1/\mathsf{poly}(n)]$ which itself is inside $1\mathsf{RE}[\mathsf{negl}(n), \mathsf{negl}(n)]$ by Key Lemma. Finally, note that $1\mathsf{RE}[\mathsf{negl}(n), \mathsf{negl}(n)]$ is equivalent to NISZK up to some polynomial-size advice [7, Therorem 1]. This completes the proof.

Proof sketch of Theorem 3 (from SZK to OWFs via error removal of 1RE). Let $\Pi \in 1\mathsf{RE}[d(n), 0]$ be a worst-case-hard problem with encoding Enc and one-sided simulator Sim.[1] The circuits Enc and Sim are randomized and take random coins as input. Without loss of generality, we can assume that they both use κ bits of randomness. Note that $\kappa = \omega(\log n)$ since otherwise $\Pi \in \mathsf{BPP}$. Let $C_0(\cdot)$ and $C_1[y](\cdot)$ respectively denote $\mathsf{Sim}(\cdot)$ and $\mathsf{Enc}(y;\cdot)$ whose randomness is considered as the input. Let $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\} \to \{0,1\}^*$ be the following function:

$$\mathsf{F}(b, r) := \begin{cases} C_0(r) & b = 0 \\ C_1[x](r) & b = 1 \end{cases}$$

where x is a fixed YES instance of Π. We show that F is a weak OWF, meaning that the probability of any polynomial-time adversary successfully inverting it on a random κ-bit input is at most $3/4$. This implies the existence of OWFs via Yao's hardness amplification [65]. Assume that there exists an inverter $\mathcal{A}$ for F that succeeds with probability larger than $3/4$. We use $\mathcal{A}$ to decide any given instance y of Π as follows: Compute $C_1(r)$, and $C_1[y](r)$ for a random value r, and then return $C_b(r)$ to $\mathcal{A}$. If $b = 0$, then $\mathcal{A}$ receives an instance of the function F and can therefore invert it. This does not help us with solving Π. Conversely, when $b = 1$, if y is a YES instance, then by the privacy of the encoding, the statistical distance between $C_1[y]$ and $C_1[x]$ would be at most $2d$. Therefore, $\mathcal{A}$ would succeed to invert it with probability $> 3/4 - 2d$. On the other hand, when $b = 1$, if y is a NO instance, then $C_1[y]$ and $C_1[x]$ have distinct support. This is because the encoding has perfect correctness. Therefore, if $b = 1$ and y is a NO instance, then there would be no preimage for the value that $\mathcal{A}$ aims to invert, which results in $\mathcal{A}$ failing. Now, assume that $d \leq 1/4$. Then we have $3/4 - 2d \geq 1/4$ and, therefore, there is a constant gap between the success rate of $\mathcal{A}$ for YES and NO instances of y. By repeating this test for several times, we can decide y by observing the success rate of $\mathcal{A}$. Therefore, F is a weak OWF in this case.

So far, we have required that $d \leq 1/4$. Fortunately, this limitation can be lifted according to Key Lemma since $1\mathsf{RE}[1 - 1/\mathsf{poly}(n), 0] \subseteq 1\mathsf{RE}[1/4, 0]$. Therefore the worst-case hardness of $1\mathsf{RE}[1 - 1/\mathsf{poly}(n), 0]$ implies OWFs. On the other hand, note that $1\mathsf{RE}[\mathsf{negl}(n), \mathsf{negl}(n)] = \mathsf{NISZK}$ [7] (up to non-uniform advice) and that NISZK is worst-case hard if and only if SZK is worst-case hard [34]. Hence, the worst-case hardness of SZK implies the worst-case hardness of $1\mathsf{RE}[\mathsf{negl}(n), \mathsf{negl}(n)]$ and if $1\mathsf{RE}[\mathsf{negl}(n), \mathsf{negl}(n)] \subseteq 1\mathsf{RE}[1/2 - 1/\mathsf{poly}(n), 0]$, then OWFs exist.

Proof sketch of Theorem 4 (from weak & imperfect iO to OWFs). To obtain the theorem, we analyze the statistical behavior of obfuscation. An algorithm is said to be $\alpha(n)$-statistical obfuscation (sO) if it satisfies the following indistinguishability property: the obfuscations of any two functionally-equivalent circuits on n-bit inputs are at most $\alpha(n)$-statistically far. We show that if sO with distance $\alpha(n) = 1 - 1/\mathsf{poly}(n)$ and error

[1] Note that since $\Pi \in 1\mathsf{RE}$, the simulator Sim works only on YES instances. We therefore omit $\chi_\Pi(x)$ from the input of Sim, as it is always 1 in this case.

$\varepsilon(n) = 1/2 - 1/\mathsf{poly}(n)$ exists, then $\mathsf{NP} \subseteq \mathsf{coNP/Poly}$, i.e., the Polynomial Hierarchy collapses to its third level.[1]

The key idea is to leverage sO to build a randomized encoding for the unique satisfiability (UniqueSat) problem. This problem asks to decide whether a 3CNF formula over n variables has a satisfiable assignment under the promise that it has at most one of such. Works of [38, 44] observe that sO gives a worst-case to average-case reduction for UniqueSat. We follow the same blueprint. More precisely, if a circuit C is a YES instance of UniqueSat, i.e., it evaluates to 1 on exactly one input, then a random shift of C, namely the circuit $C_z(x) = C(x \oplus z)$, where $z \xleftarrow{\$} \{0,1\}^n$, has a truth-table whose distribution is identical to that of a random point function. Therefore, $\mathrm{sO}(C_z)$ is statistically close to $\mathrm{sO}(\mathbb{1}_{x=x^*})$, where $x^* \xleftarrow{\$} \{0,1\}^n$. This gives us a one-sided randomized encoding for UniqueSat, with encoding $\mathsf{Enc}(C; r, z) = \mathrm{sO}(C_z; r)$ and simulator $\mathsf{Sim}(1; r, z) = \mathrm{sO}(\mathbb{1}_{x=z}; r)$ (which works for YES instances). As a result $\mathsf{UP} \subseteq 1\mathsf{RE}[\alpha, \varepsilon]$, where α is the statistical indistinguishability bound of sO and ε is the error.[2] By using Key Lemma, we can amplify the privacy and the correctness which leads to the following:

$$\mathsf{UP} \subseteq 1\mathsf{RE}[\alpha, \varepsilon] \subseteq 1\mathsf{RE}\left[\mathsf{negl}(n), \mathsf{negl}(n)\right]/\mathsf{poly} = \mathsf{NISZK/poly}\,.$$

Hence, UP collapses to SZK/poly. Finally, via the Valiant-Vazirani polynomial-time reduction of 3Sat to UniqueSat, this sequence implies that $\mathrm{3Sat} \in \mathsf{SZK/poly}$, meaning that the Polynomial Hierarchy collapses to its third level.

In fact, a two-sided randomized encoding for UniqueSat can be obtained by the exact same idea of leveraging the obfuscation. This is because for a non-satisfiable circuit C, we have that $\mathrm{sO}(C; r) \approx \mathrm{sO}(Z; r)$ where Z is the zero function. However, the rest of the argument does not hold due to lack of a privacy amplification and error reduction for two-sided encodings.

We now draw the final conclusion of this part. The preceding analysis implies that, unless the PH collapses, any (α,ε)-sO requires either α to be negligibly close to 1, or ε to be negligibly close to 1/2. Now consider an iO with computational distance $\tau(n) = 1 - 1/\mathsf{poly}(n)$ and error $\varepsilon(n) = 1/2 - 1/\mathsf{poly}(n)$. It follows that for any such iO, there exist two families of circuits $\{C_n\}_n$ and $\{C'_n\}_n$ such that their obfuscated variants are statistically far, as the statistical distance has to be $\alpha(n) = 1 - \mathsf{negl}(n)$ (or else, the PH collapses), and computationally close, as we assumed the computational distance to be $\tau(n) = 1 - 1/\mathsf{poly}(n)$. This statistically-far-but-computationally-close pair of distributions is an *EFI pair*, which is known to imply OWFs [12, 33, 51]. Putting everything together, assuming the existence of iO with computational distance $\tau(n) = 1 - 1/\mathsf{poly}(n)$ and error $\varepsilon(n) = 1/2 - 1/\mathsf{poly}(n)$ implies the existence of OWFs unless the PH collapses to its third level.

1.3 Other Related Works

Hirahara and Nanashima [37] study the existence of OWFs by leveraging Kolmogrov complexity tools. They show that if NP is polynomial-time worst-case hard and moreover $\mathsf{NP} \subseteq \mathsf{CZK}$ then OWFs exist. The complexity class CZK is the computational variant of SZK where the zero-knowledge simulator offers only computational indistinguishability guarantees. In fact the condition is equivalent to the polynomial-time worst-case hardness of CZK and the membership of the decision variant of the minimum circuit size problem to this class. [1]

Chakraborty, Hulett, and Khurana [22] show how to build a OWF from the worst-case hardness of NP if NP is contained in NICZK (computational zero-knowledge) with $\varepsilon_{zk} + \sqrt{\varepsilon_s} <$

[1] Recall that an obfuscation O has error $\varepsilon(n)$ if for every circuit over n-bit inputs, it holds that $\Pr[\exists x \in \{0,1\}^n : \mathsf{O}(C)(x) \neq C(x)] \leq \varepsilon(n)$.

[2] This part must be done carefully and requires a change of variable from n, the size of the input, to $|C|$ the instance size since the class 1RE is parameterized by the instance size. Also note that UniqueSat is complete for UP.

[1] The search variant of this problem roughly asks to find the minimum-size circuit that outputs the truth-table of a given function.

1, where ε_{zk} and ε_s are respectively zero-knowledge and soundness errors. Their result is based on the work of Hirahara and Nanashima [37].

IV-OWSGs seem to be weaker than OWFs: there are (quantum or classical) oracles relative to which IV-OWSGs exist while OWFs do not (see [45, 46]). Khurana and Tomer [43] show the existence of IV-OWSGs with mixed outcome based on the worst-case hardness of $\mathsf{P}^{\#\mathsf{P}}$ and an extra ad-hoc assumption. More generally, it is known that IV-OWSGs with mixed states are equivalent to EFIs and therefore the existence of the latter implies the existence of the former [42]. Although pure IV-OWSGs imply EFIs [42], the other direction is not known yet. Bostanci et al. [16] construct IV-OWSG with pure outcome via the worst-case hardness of the so-called Hamiltonian Phase State Problem.

1.4 Open Questions

The most important question that remains open is an error removal for $\mathsf{1RE}$. As we showed, such an amplification would immediately implies OWFs from the worst-case hardness of SZK. Even a weaker amplification such as the following one yields the same result: if we can efficiently (up to some advice) recognize the instances or/and the random coins on which the encoding errs. Also working with the sparsely lossy characterization of $\mathsf{1RE}$ might be helpful. We leave this question to the future work.

2 Preliminaries

In this work, we always consider non-uniform algorithms. All classical algorithms are quantum algorithms, therefore, we mostly use the quantum formalism for generalization and simplification. When the distinction is necessary, we explicitly mention it in the beginning of a section or inside a statement.

Notation. We let n denote the security parameter, and all variables are implicitly parameterized by n. We let MS_n denote the set of all mixed states over n qubits and we define $\mathsf{MS}_* := \cup_{n=1}^{\infty}\mathsf{MS}_n$. For a positive integer n, we let $[n]$ denote $\{1, 2, \cdots, n\}$. The set of all permutations over $[n]$ is $\mathfrak{S}_n$. We abuse the notation and use the same symbol to refer to the uniform distribution over all permutations of $[n]$. The set of natural numbers $\{1, 2, 3, \cdots\}$ is denoted by $\mathbb{N}$. We denote by $\mathbb{R}^+$ the set of positive real numbers.
A function $f(n)$ is $\mathsf{poly}_\ell(n)$ if $f(n) = O(n^\ell)$ (we drop ℓ when the degree is not specified), is $\mathsf{negl}(n)$ if $f(n) = 1/n^{\omega(1)}$, and $\mathsf{subexp}(n)$ if $f(n) = 2^{O(n^c)}$ for some constant $c < 1$.

Uniform and s-Uniform Distributions. For any set S, we let $\mathcal{U}_S$ denote the uniform distribution over S. A distribution is called s-uniform if it is uniform over a multiset of s elements.

Promise Problems. A Promise Problem Π consists of two disjoint sets Π_Y and Π_N in $\{0,1\}^*$, respectively referred to as the set of YES and NO instances. Problem Π asks to decide whether a given instance, which is promised to lie in $\Pi_Y \cup \Pi_N$, belongs Π_Y or Π_N.

Definition 1 (Characteristic Function of a Promise Problem). *For a promise problem Π, its characteristic function is the map $\chi_\Pi(x) : \{0,1\}^* \to \{0, 1, \star\}$ given by*

$$\chi_\Pi(x) = \begin{cases} 1 & \text{if } x \in \Pi_Y \\ 0 & \text{if } x \in \Pi_N \\ \star & \text{otherwise} \end{cases}.$$

In this work, we consider worst-case hardness as introduced below.

Definition 2. *A promise problem Π is worst-case hard, if for any (randomized or quantum depending on the context) polynomial-time algorithm $\mathcal{A}$ and any sufficiently large $n \in \mathbb{N}$, there exists an input $x \in (\Pi_Y \cup \Pi_N) \cap \{0,1\}^n$ such that $\Pr[\mathcal{A}(x) = \chi_\Pi(x)] < 2/3$.*

Classical information. Given two probability distributions X and Y over Σ, their statistical distance, also called total variation distance, is defined as

$$\Delta(X,Y) := \frac{1}{2}\sum_{x\in\Sigma} |\Pr(X=x) - \Pr(Y=x)| \ .$$

The Kullback–Leibler divergence or classical relative entropy of X with respect to Y is defined as

$$D_{\mathrm{KL}}(X\|Y) := \sum_{x\in\Sigma} \Pr(X=x)\log\left(\frac{\Pr(X=x)}{\Pr(Y=x)}\right) \ .$$

Lemma 2.1 (Chernoff-Hoeffding Bound). *Let $X_1, X_2, \cdots, X_k$ be mutually independent random variables in $[0,1]$ and $\mu := \mathbb{E}[\sum_{i=1}^k X_i]$. For every $t > 0$, it holds that $\Pr[|\sum_{i=1}^k X_i - \mu| > t] \le 2\exp(-2t^2/k)$.*

Quantum information. For any two mixed states ρ and σ, we let $D_{\mathrm{tr}} := \mathrm{Tr}(|\rho-\sigma|)/2$ denote their trace distance. Let $R : \{0,1\}^n \to \mathsf{MS}_m$ be any quantum mapping and X a random variable supported over $\{0,1\}^n$. We let

$$\rho_{X,R(X)} := \sum_{x\in\{0,1\}^n} \Pr_X(x)\,|x\rangle\langle x| \otimes R(x) \ . \tag{5}$$

For a mixed state ρ, we let $S(\rho) := \mathrm{Tr}(\rho\log_2\rho)$ denote the Von Neumann entropy of ρ. The quantum mutual information of two subsystems A and B is defined as follows. Let ρ_{AB} be their joint state, then

$$I_q(A;B)_\rho := S(\rho_A) + S(\rho_B) - S(\rho_{AB}) \ ,$$

where $\rho_A = \mathrm{Tr}_B(\rho_A B)$ and $\rho_B = \mathrm{Tr}_A(\rho_A B)$. For the sake of simplicity, we sometimes drop the subscripts q and ρ in I_q. When working with quantum systems A, B, the notation $I(A;B)$ implicitly refers to $I_q(A;B)$.

The following lemma states that if the outcome of a measurement is close to deterministic, then it must not alter much the state.

Lemma 2.2 (Gentle Measurement Lemma [64]). *Let ρ be a mixed state and $\{\Lambda, I-\Lambda\}$ a two-outcome POVM with $\mathrm{Tr}(\Lambda\rho) \ge 1-\varepsilon$, then $D_{tr}(\rho,\rho') \le \sqrt{\varepsilon}/2$, where $\rho' = \frac{\sqrt{\Lambda}\rho\sqrt{\Lambda}}{\mathrm{Tr}(\Lambda\rho)}$.*

For two quantum states σ, ρ stored in two different registers A, B, the swap test is executed on the registers A, B and a control register C initialized to $|1\rangle\langle 1|$. It applies Hadamard on C, swaps A and B conditioned on C, and measures B on the Hadamard basis.

Lemma 2.3 (SWAP Test [20]). *The SWAP test on input (σ,ρ) outputs 1 with probability $(1+\mathrm{Tr}(\rho\sigma))/2$, in which case we say that it passes the test. For pure states $|\sigma\rangle, |\rho\rangle$, it equals to $(1+|\langle\rho|\sigma\rangle|^2)/2$.*

Given that the trace distance of two pure states $|\sigma\rangle, |\rho\rangle$ can be expressed in terms of the inner product uniquely as $\sqrt{1-|\langle\rho|\sigma\rangle|^2}$, the SWAP test can also be used to calculate their trace distance.

Distance notation. To avoid repeated proofs, we use the following definition to unify the distance analysis in both classical and quantum settings.

Definition 3. *We use the notation $\|X - Y\|$ to refer to the statistical distance $\Delta(X, Y)$ when variables X, Y are classical distributions, or the trace distance when they are quantum states.*

We will use the following lemmas.

Lemma 2.4 (Direct Product Lemma [56, 61]). *Let X, Y be two quantum states or classical distributions such that $\|X - Y\| = \varepsilon$. For any k, it holds that*

$$1 - 2\exp\left(\frac{-k\varepsilon^2}{2}\right) \leq \left\|X^{\otimes k} - Y^{\otimes k}\right\| \leq k\varepsilon\,.$$

Lemma 2.5 (XOR Lemma [56,61]). *Let XOR be a function that maps the tuple $(X_0, X_1, 1^k)$ to (Z_0, Z_1) as follows:*

Z_0 : Uniformly sample $(b_1, \cdots, b_k) \in \{0,1\}^k$ such that $b_1 \oplus \cdots \oplus b_k = 0$ and output a sample from $X_{b_1} \otimes \cdots \otimes X_{b_k}$.

Z_1 : Uniformly sample $(b_1, \cdots, b_k) \in \{0,1\}^k$ such that $b_1 \oplus \cdots \oplus b_k = 1$ and output a sample from $X_{b_1} \otimes \cdots \otimes X_{b_k}$.

Then $\|Z_0 - Z_1\| = \|X_0 - X_1\|^k$.

Lemma 2.6. *Let (X_0, X_1) and (Y_0, Y_1) be such that $\|X_0 - Y_0\| \geq \varepsilon$ and $\|X_1 - Y_1\| \geq \delta$. Let $(P, \cdot) := \mathsf{XOR}(X_0, X_1, 1^k)$ and $(Q, \cdot) := \mathsf{XOR}(Y_0, Y_1, 1^k)$ for some $k \geq 2$. Then $\|P - Q\| \geq (\varepsilon + \delta)/2$.*

Proof. Consider the first component of P. The marginal distribution of b_1 (see Lemma 2.5) is uniform. Thus with half probability $b_1 = 0$ and the first component is selected from X_0. Same holds for Q: when $b_1 = 0$, its first component is Y_0. Note that the distance is bounded by ε in this case. Same holds for $b_1 = 1$ with the lower bound δ. This completes the proof. □

Permutation-invariant function. A function is permutation-invariant if applying any permutation to its inputs does not alter its output value. This includes a wide range of functions, such as And, Or, Maj, and Parity. We use the following technical lemma about non-constant permutation-invariant functions.

Lemma 2.7. *Let $f : \{0,1\}^m \to \{0,1\}$ be a non-constant permutation-invariant function. Then there exists an integer $1 \leq p \leq m$ such that*

$$f(\underbrace{1, 1, \ldots, 1}_{p-1}, 0, 0, \ldots, 0) = 0\,, \quad \text{and} \quad f(\underbrace{1, 1, \ldots, 1}_{p}, 0, 0, \ldots, 0) = 1\,.$$

We let $p(f)$ denote the minimum choice of such an integer.

Proof. The set $\{0,1\}^m$ can be partitioned into $m + 1$ equivalence classes where each class consists of strings with the same number of 1's. We note that the result of a permutation on an input falls in the same equivalence class. Therefore, since the function is permutation-invariant, then the evaluation of f over each input is determined by its class. Because the function is non-constant, there must exist two consecutive classes (the classes can be ordered by the number of 1's that they represent) with different evaluation under f. This completes the proof. □

Complexity class SZK. There are various equivalent definitions of the complexity class SZK. We start with the standard one.

Definition 4 (Class SZK). *For any $c, s, \varepsilon : \mathbb{N} \to [0,1]$ we let $\mathsf{SZK}[c(n), s(n), \varepsilon(n)]$ be the set of all promise problems Π that admit a public coin interactive proof system $(\mathcal{P}, \mathcal{V})$ such that:*

- ***Completeness**: For any $x \in \Pi_Y \cap \{0,1\}^n$, it holds that $\Pr[\langle \mathcal{P}, \mathcal{V} \rangle(x) = 1] \geq 1 - c(n)$.*
- ***Soundness**: For any $x \in \Pi_N \cap \{0,1\}^n$ and every prover $\mathcal{P}^*$, it holds that $\Pr[\langle \mathcal{P}^*, \mathcal{V} \rangle(x) = 1] \leq s(n)$.*
- ***Zero-Knowledge**: There exists a polynomial-time algorithm Sim such that for every $x \in \Pi_Y \cap \{0,1\}^n$, it holds that*

$$\|\mathsf{Sim}(x) - \mathrm{View}_{\mathcal{P},\mathcal{V}}(x)\| \leq \varepsilon(n)\,.$$

We let $\mathsf{SZK} := \mathsf{SZK}[\mathsf{negl}(n), \mathsf{negl}(n), \mathsf{negl}(n)]$.

For certain parameters, the following problem is complete for SZK.

Definition 5 (Statistical Difference). *Let $\alpha, \beta \in [0,1]$ such that $\alpha < \beta$. Given two circuits $\mathcal{C}_0$ and $\mathcal{C}_1$, with the promise that either $\|C_0 - C_1\| \geq \beta$ (corresponds to NO instances) or $\|C_0 - C_1\| \leq \alpha$ (corresponds to YES instances). The $\mathrm{SD}_{\alpha,\beta}$ problem is to decide which one is the case. When it is clear from the context, we abuse the notation and let C denote its outcome distribution too.*

Sahai and Vadhan [56] show that $\mathrm{SD}_{1/3,2/3}$ is complete for SZK. They also show that SZK has a the following closure property.

Lemma 2.8 (From [56, Th. 4.9]). *Let $f : \{0,1\}^m \to \{0,1\}$ be a Boolean function. For a problem Π define $f(\Pi)$ as follows:*

$$f(\Pi)_Y := \{(x_1, \cdots, x_m) \mid f(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m)) = 1\}\,,$$

and

$$f(\Pi)_N := \{(x_1, \cdots, x_m) \mid f(\chi_\Pi(x_1), \cdots, \chi_\Pi(x_m)) = 0\}\,.$$

For every $\Pi \in \mathsf{SZK}$, it holds that $f(\Pi) \in \mathsf{SZK}$.

We also have the following lemma about SZK that is derived from [28, Thm.4.11] and [17, Thm. 4].

Lemma 2.9. *It holds that $\mathsf{SZK/poly} \subseteq \mathsf{NP/poly} \cap \mathsf{coNP/poly}$.*

Complexity class NISZK. It is the class of problems that have non-interactive zero-knowledge proofs.

Definition 6 (Class NISZK). *It is defined similarly to $\mathsf{SZK}[c(n), s(n), \varepsilon(n)]$ with the following differences: (1) Only one single message from the prover $\mathcal{P}$ to the verifier $\mathcal{V}$ is communicated. (2) There exists a polynomial-time algorithm $\mathsf{Deal}(1^n)$ that is run before the communication to obtain a uniformly random string σ. The parameter σ is given to both $\mathcal{P}$ and $\mathcal{V}$ and is considered as part of the view of the protocol. We also let $\mathsf{NISZK} := \mathsf{NISZK}[\mathsf{negl}(n), \mathsf{negl}(n), \mathsf{negl}(n)]$.*

The worst-case hardness of NISZK and SZK are equivalent.

Lemma 2.10 (Adapted from [34, Th. 1.5]). *It holds that $\mathsf{NISZK} \subseteq \mathsf{P/poly}$ if and only if $\mathsf{SZK} \subseteq \mathsf{P/poly}$.*

The randomized encodings are inside NISZK and SZK.

Lemma 2.11 (From [7]). *It hold that $\mathsf{RE} \subseteq \mathsf{SZK}$, $\mathsf{1RE} \subseteq \mathsf{NISZK}$, and $\mathsf{NISZK} \subseteq \mathsf{1RE/poly}$.*

Polynomial Hierarchy. We let $\Sigma_0^p := \Pi_0^p := \mathsf{P}$ be the class of problems solvable in polynomial time. The levels of the Polynomial Hierarchy PH are defined as follows: for every $i \in \mathbb{N}$, we let

$$\Sigma_{i+1}^p := \mathsf{NP}^{\Sigma_i^p} \text{ and } \Pi_{i+1}^p := \mathsf{coNP}^{\Sigma_i^p}.$$

It is a common belief that PH does not collapse to any of its levels, namely, it is unlikely to have $\mathsf{PH} = \Sigma_i^p = \Pi_i^p$ for any $i \geq 1$. The following result will be used in our theorems.

Lemma 2.12 (From [66]). *If* $\mathsf{NP} \subseteq \mathsf{coNP/poly}$, *then* $\mathsf{PH} = \Sigma_3^p = \Pi_3^p$.

Cryptographic primitives. One-way functions are defined as follows:

Definition 7 (One-Way Functions). *A family of non-uniform PPT algorithms* $\mathsf{F} := \{\mathsf{F}_n\}_{n\in\mathbb{N}}$ *is said to be one-way (OWF) if for all sufficiently large* n *and any polynomial-time algorithm* $\mathcal{A}$, *it holds that*

$$\Pr_{x\sim\mathcal{U}_{\{0,1\}^n}}[\mathsf{F}(\mathcal{A}(\mathsf{F}(x))) = \mathsf{F}(x)] \leq \mathsf{negl}(n).$$

Furthermore, we say that F *is weak-one-way if the above probability is bounded by* $1 - 1/\mathsf{poly}(n)$.

It is shown by [65] that weak one-way functions imply one-way functions.
In this work, we consider the inefficient-verifier one-way state generators with pure output.

Definition 8 (One-Way State Generators). *A one-way state generator (OWSG) is a tuple of algorithms* $\mathsf{G} := (\mathsf{KeyGen}, \mathsf{StateGen}, \mathsf{Ver})$ *with the following specification:*

- $\mathsf{KeyGen}_h(1^n) \to k$*: is a QPT algorithm that given the security parameter* 1^n *and a* $\mathsf{poly}(n)$*-size classical advice* h, *outputs a classical string* $k \in \{0,1\}^n$*;*
- $\mathsf{StateGen}(k) \to |\psi_k\rangle$*: is a QPT algorithm that given a classical string* k, *outputs an* m*-qubit quantum state;*
- $\mathsf{Ver}(k, |\psi\rangle) \in \{0,1\}$*: is a (possibly unbounded) algorithm that given a classical string* k *and a quantum state* $|\psi\rangle$ *outputs either* 0 *or* 1.

Further, they satisfy the following properties:

1. ***Correctness.*** *Outputs of the samplers* $(\mathsf{KeyGen}, \mathsf{StateGen})$ *pass the verification with overwhelming probability, i.e.,*

$$\Pr_{\substack{k\leftarrow\mathsf{KeyGen}_h\\ |\psi_k\rangle\leftarrow\mathsf{StateGen}(k)}}[\mathsf{Ver}(k, |\psi_k\rangle) = 1] \geq 1 - \mathsf{negl}(n).$$

2. ***Security.*** *For every non-uniform* $T(n)$*-time adversary* $\mathcal{A}$, *and any polynomial* $t(n)$

$$\Pr_{\substack{k\leftarrow\mathsf{KeyGen}_h\\ |\psi_k\rangle\leftarrow\mathsf{StateGen}(k)\\ k'\leftarrow\mathcal{A}(|\psi_k\rangle^{\otimes t};h)}}\left[\mathsf{Ver}(k', |\psi_k\rangle) = 1\right] \leq \mathsf{negl}(n).$$

Furthermore, we say that G *is a weak-OWSG if the inequality concerning security (Property 2) is bounded by* $1 - 1/\mathsf{poly}(n)$.

It is shown in [50] that weak OWSGs imply OWSGs.

3 Complexity Classes

In this section, we define useful complexity classes. We begin by the class of problems that have "lossy" reductions. In our analysis, we consider both randomized functions and quantum mappings. All the statements hold with respect to both cases. We explicitly highlight the distinction when the analysis requires to distinguish between the two cases.

Classically, a randomized function $R : \{0,1\}^* \to \{0,1\}^*$ is said to be λ-lossy for a distribution X if $I(X; R(X)) \leq \lambda$. Below, we also consider general mappings with classical input and quantum output.

Definition 9 (Lossy Mapping on a Fixed Distribution). *Let $\lambda, m \geq 0$. Let $R : \{0,1\}^* \to S$ be a mapping, where $S = \{0,1\}^*$ (classical mapping) or $S = \mathsf{MS}_*$ (quantum mapping). We say that R is λ-lossy for an m-tuple distribution $X = (X_1, X_2, \ldots, X_m)$ over $\{0,1\}^*$, if it holds that*

$$I(X; R(X)) \leq m\lambda \ .$$

For the sake of simplicity, we say that R is λ-lossy, if it is λ-lossy for all m-tuple distributions.

As we will see later, it is useful to limit the input distributions to have a particular "splitting" structure.

Definition 10 (Splitting Distribution). *Let $S_0, S_1 \subseteq \{0,1\}^*$ be two disjoint sets. We say that a distribution $X = (X_1, \ldots, X_m)$ splits over the pair (S_0, S_1) if for each $i \in [m]$, either $\mathrm{Supp}(X_i) \subseteq S_0$ or $\mathrm{Supp}(X_i) \subseteq S_1$.*

An example of such distributions is (X, Y) where X and Y have disjoint supports.

In this work, we consider distributions that are not only sparse but also are uniform.

Definition 11 (Pairwise Independent B-Uniform Distribution). *We say that a distribution $X = (X_1, \ldots, X_m)$ is pairwise-independent B-uniform if X_i's are independent and each of them follows a (not necessarily similar) uniform distribution over a multiset of size at most B.*

For example, distribution $(\mathcal{U}_n, \cdots, \mathcal{U}_n)$ is pairwise independent 2^n-uniform.

For the sake of generality, we consider the class of problems that admit lossy f-reductions. Recall f reductions below.

Definition 12 (f-Reduction). *Let n, m be positive integers, and $\mu : \mathbb{N} \to [0, 1/2]$ be a function of n. Let $f : \{0,1\}^m \to \{0,1\}$, and Π be a promise problem. An f-reduction for Π with error μ is a mapping $R : \{0,1\}^{mn} \to S$, where $S = \{0,1\}^*$ (classical) or $S = \mathsf{MS}_*$ (quantum), for which there exists an unbounded algorithm $\mathcal{M}$, such that for all $(x_1, \ldots, x_m) \in ((\Pi_Y \cup \Pi_N) \cap \{0,1\}^n)^m$,*

$$\Pr\left[\mathcal{M}(R(x_1, \ldots, x_m)) = f\left(\chi_\Pi(x_1), \ldots, \chi_\Pi(x_m)\right)\right] \geq 1 - \mu(n)\,,$$

where the probability is taken over the randomness of R and $\mathcal{M}$. [1]

Note that reductions are classically defined from a problem Π to another problem Σ such that they map YES (and NO) instances of Π to YES (and NO) instances of Σ. Any (possibly unbounded) solver for Σ can play the role of $\mathcal{M}$ in our definition.

We introduce the class of *sparsely lossy problems* below.

Definition 13 (Class (Q)SL). *Let n be the instance size. Let $m \in \mathbb{N}$, $\lambda \in \mathbb{R}^+$, and $\mu \in [0, 1/2]$ be all functions of n and $f : \{0,1\}^m \to \{0,1\}$ be a Boolean function. The class $\mathsf{SL}_f[\lambda, \mu]$ consists of all promise problems Π that has a polynomial-time f-reduction with error μ with the following property:*

[1] When considering quantum mappings, $\mathcal{M}$ can be a binary quantum measurement.

- *For every polynomial $B(n)$ and for all sufficiently large n, reduction R is $\lambda(n)$-lossy (per Definition 9) for all distributions $(X_1, \cdots, X_m)$ that split over $(\Pi_N \cap \{0,1\}^n, \Pi_Y \cap \{0,1\}^n)$ (per Definition 10) and are pairwise independent B-uniform (per Definition 11).*

The class QSL_f is defined similarly by replacing classical reductions with quantum reductions. When f is the identity function, we remove the dependency on f from the notation.

Roughly speaking, this class contains those problems with reductions that lose information on sparse uniform distributions that split over YES and NO instances.

Remark 2 (One-sided lossiness). We also consider a specific type of SL classes where the lossiness is only guaranteed for YES instances. We let $1\mathsf{SL}[\lambda(n)]$ be defined as above for f being the identity function where the reduction is $\lambda(n)$-lossy for all distributions X over $\Pi_Y \cap \{0,1\}^n$ and are B-uniform. It trivially includes $\mathsf{SL}[\lambda(n)]$. The class $1\mathsf{QSL}$ is defined similarly.

We also recall the class of randomized encodings.

Definition 14 (Class (Q)RE (from [1,7])). *Let n be the instance size. Let $d \in [0,1]$, and $\mu \in [0,1/2]$ be both functions of n. The class $\mathsf{RE}[d,\mu]$ consists of all promise problems Π for which there exists a polynomial-time algorithm Enc that has the following properties:*

- *(μ-**Correctness**) There exists an algorithm Dec such that for all $x \in \{0,1\}^n$:*

$$\Pr\left[\mathsf{Dec}(\mathsf{Enc}(x)) \neq \chi_\Pi(x)\right] \leq \mu(n)\,.$$

- *(d-**Privacy**) There exists a polynomial-time algorithm Sim such that for all $x \in \{0,1\}^n$:*

$$\|\mathsf{Sim}(1^n, \chi_\Pi(x)) - \mathsf{Enc}(x)\| \leq d(n)\ ,$$

as distributions. For simplicity, we drop the argument 1^n, when input size is clear from the context.

The class $\mathsf{QRE}[d,\mu]$ is defined similarly by replacing classical algorithm Enc with quantum algorithms. We let RE and QRE denote the variants where both d and μ are negligible.

The following inclusions are trivial.

Proposition 1. *It hold that $\mathsf{BPP} \subseteq \mathsf{RE}[0, \mathsf{negl}(n)]$ and $\mathsf{BQP} \subseteq \mathsf{QRE}[0, \mathsf{negl}(n)]$.*

Remark 3 (One-sided encodings). The one-sided variant of randomized encodings are defined similarly except that the privacy is only guaranteed for YES instances and not all instances. More precisely, there exists a polynomial-time algorithm Sim that is statistically close to $\mathsf{Enc}(x)$ for every YES instance x. We let $1\mathsf{RE}[d,\mu]$ and $1\mathsf{QRE}[d,\mu]$ denote these classes in the classical and quantum settings, respectively. The one-sided variant is possibly a larger class. AppleBaum and Raykov [7] show that $1\mathsf{RE}[\mathsf{negl}(n), \mathsf{negl}(n)]$ is equal to the class of problems that have non-interactive statistical zero-knowledge proofs (NISZK).

Remark 4 (Pure-outcome encodings). In the above definition of $(1)\mathsf{QRE}$, there is no restriction on the type of the output of Enc; it can be a mixed or a pure state. We also define the more restrictive class $(1)\mathsf{QRE}_{\mathsf{pure}}$ to be the class of those quantum encodings whose $\mathsf{Enc}(\cdot; r)$ function returns a pure state for any fixed classical randomness r (note that the expectation $\mathsf{Enc}(\cdot; r)$ over the choice of r is not pure).

4 Sparsely Lossy Reductions Are Randomized Encodings

In this section, we show that SL_f is included in RE up to some classical advice for a variety of Boolean functions f as below:

Definition 15 (Permutation-Invariant Function). *A Boolean function $f : \{0,1\}^m \to \{0,1\}$ is said to be permutation-invariant if for every $\pi \in \mathfrak{S}_m$, it holds that $f(\pi(b_1, b_2, \ldots, b_m)) = f(b_1, b_2, \ldots, b_m)$.*

The functions And, Or, and Maj that were considered in [9, 28, 30] are all permutation-invariant. Moreover, the functions Parity and $\textsc{Mod}_k$ are of this type as well as $\textsc{Threshold}_k$.

Main theorem of this section is the following:

Theorem 4.1. *For any permutation-invariant function f and any polynomial $q(n)$, it hold that $\mathsf{SL}_f[\lambda,\mu] \subseteq \mathsf{RE}[\delta + 1/q, \mu]/\mathsf{poly}$ and $\mathsf{QSL}_f[\lambda,\mu] \subseteq \mathsf{QRE}[\delta + 1/q, \mu]/\mathsf{poly}$ where $\delta := \min\{\sqrt{(\lambda \ln 2)/2}, 1 - 2^{-\lambda-2}\}$. The advice size itself is a polynomial in q.*

The technical core for this theorem is an extended variant of *disguising lemma*. The disguising lemma first appeared in [28] and was improved by [9]. Those variants rely on lossiness with respect to all possible input distributions. Instead, we show that the lemma remains almost intact for maps that are lossy only over splitting uniform distributions with sparse support. This is obtained by a more refined analysis but yet very similar to those of [9, 28]. The proof is deferred to Appendix A.

Lemma 4.1 (Extended Disguising Lemma). *Let n, m, m_0, m_1 be positive integers such that $m = m_0 + m_1 + 1$, and $R : \{0,1\}^* \to \mathsf{MS}_*$ be any m-variate quantum mapping. Further, let $S_0, S_1 \subseteq \{0,1\}^n$ be two disjoint sets, d be a positive integer, $\varepsilon > 0$ be real, and $s := \lceil n \ln 2/(2\varepsilon^2) \rceil$.*

For any choice of positive real ℓ, if R is ℓ-lossy for all (ds)-uniform distributions that split over the pair (S_0, S_1), then there exist two collections $K_1, \ldots, K_s$ and $T_1, \ldots, T_s$ of multisets of d elements respectively contained in S_0 and S_1, such that

1. *for any $y \in S_0$, it holds that*

$$\mathop{\mathbb{E}}_{\substack{a \sim \mathcal{U}_{[s]} \\ \pi \sim \mathfrak{S}_m}} \left[\left\| R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m_0}, y, \mathcal{U}_{T_a}^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes (m_0+1)}, \mathcal{U}_{T_a}^{\otimes m_1}\right)\right) \right\| \right] \leq \delta + \frac{2(m+1)}{d+1} + 2\varepsilon \ ,$$

2. *and for any $y \in S_1$, it holds that*

$$\mathop{\mathbb{E}}_{\substack{a \sim \mathcal{U}_{[s]} \\ \pi \sim \mathfrak{S}_m}} \left[\left\| R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m_0}, y, \mathcal{U}_{T_a}^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m_0}, \mathcal{U}_{T_a}^{\otimes (m_1+1)}\right)\right) \right\| \right] \leq \delta + \frac{2(m+1)}{d+1} + 2\varepsilon \ ,$$

where

$$\delta := \min\left\{ \sqrt{\frac{\ell \ln 2}{2m}}, 1 - 2^{-\frac{\ell}{m} - 2} \right\}.$$

Note that the states inside the trace distance are mixed states since the inputs of R are randomized classical distributions.

Remark 5. If the mapping R is 1-variate and lossy only on S_0-supported (or S_1-supported) distributions, then only Item 1 (or Item 2) holds.

Algorithm 1 Randomized Encoding for Π.

Parameters: Let

$$S_0 := \Pi_N \cap \{0,1\}^n, \quad S_1 := \Pi_Y \cap \{0,1\}^n, \quad \varepsilon := \frac{1}{4q}, \quad d := \left\lceil \frac{m+1}{\varepsilon} \right\rceil, \; s := \left\lceil \frac{n \ln 2}{2\varepsilon^2} \right\rceil \; ,$$

and $K_1, \dots K_s, T_1, \dots, T_s$ as in Lemma 4.1. ε is set such that $B = ds$ where B is defined as in Definition 13.

Input: An instance $y \in \{0,1\}^n$.

Advice: K_a, T_a, π for some uniformly chosen $a \in [s]$ and $\pi \in \mathfrak{S}_m$.

Output: A pair of circuits $(\mathsf{Sim}, \mathsf{Enc})$.

1: Let $p(f)$ be as in Lemma 2.7.

2: Let $\mathsf{Enc}(y)$ be the following circuit: it samples $\widetilde{x} \sim \left(\mathcal{U}_{K_a}^{\otimes m-p}, y, \mathcal{U}_{T_a}^{\otimes p-1}\right)$, then it outputs $R(\pi(\widetilde{x}))$.

3: Let $\mathsf{Sim}(0)$ be the following circuit: it samples $\widetilde{x} \sim \left(\mathcal{U}_{K_a}^{\otimes m-p+1}, \mathcal{U}_{T_a}^{\otimes p-1}\right)$, then it outputs $R(\pi(\widetilde{x}))$.

4: Let $\mathsf{Sim}(1)$ be the following circuit: it samples $\widetilde{x} \sim \left(\mathcal{U}_{K_a}^{\otimes m-p}, \mathcal{U}_{T_a}^{\otimes p}\right)$, then it outputs $R(\pi(\widetilde{x}))$.

Proof of Theorem 4.1. We assume that f is non-constant since otherwise the inclusion is trivial. Let $\Pi \in \mathsf{SL}_f$ or $\Pi \in \mathsf{QSL}_f$ and R be its sparsely lossy reduction. We show that $(\mathsf{Enc}, \mathsf{Sim})$ defined as in Algorithm 1 is a randomized encoding for Π. Before that, let us explain the input-output type of these circuits. Consider $X \in \{\mathsf{Sim}, \mathsf{Enc}\}$. When R is a randomized reduction, the two circuits are also randomized. Part of their randomness input is used to sample $\widetilde{x}$ and the other part is fed to R. Let κ be the size of the total randomness. For $r \in \{0,1\}^\kappa$, we let $X(r)$ denote the outcome of X given the randomness r. On the other hand, when R is quantum, the circuits will be mixed algorithms; classical randomness is required for sampling $\widetilde{x}$. Let κ' be the size of total randomness.[1] For any $r \in \{0,1\}^{\kappa'}$, we let the mixed outcome of X be $X\,|r, \mathbf{0}\rangle$ where $|\mathbf{0}\rangle$ is some appropriate-size ancilla, emphasizing its mixed classical-quantum nature. When it is not relevant, we drop the dependency on r for simplification.

The proof is exactly similar for both classical and quantum case. Recall that, we need to show correctness and privacy properties as required in Definition 14.

Correctness: It suffices to show that the distance (per Definition 3) of the outcomes $\mathsf{Enc}(y)$ and $\mathsf{Enc}(y')$ for any choice of $y \in \Pi_Y$ and $y' \in \Pi_N$ is at least $1 - 2\mu(n)$. We use the fact that the distance is equal to the advantage of the best distinguisher in the following game: Sample a uniform coin $b \sim U_{\{0,1\}}$. If $b = 0$ let $z \leftarrow \mathsf{Enc}(y')\,|r, \mathbf{0}\rangle$ where r follows the uniform distribution. Otherwise let $z \leftarrow \mathsf{Enc}(y)\,|r, \mathbf{0}\rangle$. We drop the dependency on r for simplification. On one hand, if z is computed by $\mathsf{Enc}(y')$, we have that $\widetilde{x} := (x_1, \dots, x_m) \sim \left(\mathcal{U}_{K_a}^{\otimes m-p}, y', \mathcal{U}_{T_a}^{\otimes p-1}\right)$ with $K_a \subseteq \Pi_N \cap \{0,1\}^n$ and $T_a \subseteq \Pi_Y \cap \{0,1\}^n$. Then, since $\widetilde{x}$ contains $p-1$ YES instances, by Lemma 2.7, for any $\pi \in \mathfrak{S}_m$, we have

$$f(\pi(\chi_\Pi(x_1), \dots, \chi_\Pi(x_m))) = 0 \; .$$

On the other hand, if z is computed by $\mathsf{Enc}(y)$, we have that $\widetilde{x}$ contains one more YES instance $y \in \Pi_Y \cap \{0,1\}^n$, therefore,

$$f(\pi(\chi_\Pi(x_1), \dots, \chi_\Pi(x_m))) = 1 \; .$$

[1] Note that κ and κ' are possibly different depending on how much classical randomness R requires.

Thus by the correctness of the reduction (per Definition 12), there exists an algorithm $\mathcal{M}$ such that

$$\begin{aligned}&\|\mathsf{Enc}(y')\,|\mathbf{0}\rangle - \mathsf{Enc}(y)\,|\mathbf{0}\rangle\,\|\\&\geq \mathbb{E}_{i\sim\mathcal{U}_{[m]}}\left|\Pr\left[1\leftarrow\mathcal{M}(\mathsf{Enc}(y')\,|\mathbf{0}\rangle)\right] - \Pr[1\leftarrow\mathcal{M}(\mathsf{Enc}(y)\,|\mathbf{0}\rangle)]\right|\\&\geq 1-2\mu(n)\,,\end{aligned}$$

where we used the fact that revealing π with the description of the circuits does not decrease the success probability of any distinguisher.

Privacy: Suppose $y\in\Pi_N$. We consider a modification of the distinguishing game where the random variables a and π are also given to the distinguisher. Revealing a,π along with z does not decrease the success probability of the distinguisher, thus we can bound the original distinguishing probability by the distinguishing probability of the new task. It holds that

$$\begin{aligned}&\|\mathsf{Sim}(0)\,|\mathbf{0}\rangle - \mathsf{Enc}(y)\,|\mathbf{0}\rangle\,\|\\&\qquad\leq\left\|R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p+1},\mathcal{U}_{T_a}^{\otimes p-1}\right)\right)-R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p},y,\mathcal{U}_{T_a}^{\otimes p-1}\right)\right)\right\|,\end{aligned}$$

By taking the expectation over a and π, we have

$$\begin{aligned}&\|\mathsf{Sim}(0)\,|\mathbf{0}\rangle - \mathsf{Enc}(y)\,|\mathbf{0}\rangle\,\|\\&\leq\mathbb{E}_{a\sim\mathcal{U}_{[s]},\pi\sim\mathfrak{S}_m}\left[\left\|R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p+1},\mathcal{U}_{T_a}^{\otimes p-1}\right)\right)-R\left(\pi\left(\mathcal{U}_{K_a}^{\otimes m-p},y,\mathcal{U}_{T_a}^{\otimes p-1}\right)\right)\right\|\right].\end{aligned}$$

Note that by assumption R is λ-lossy for all $\mathsf{poly}(n)$-uniform distributions that split over (Π_Y,Π_N). Therefore, one can choose $B=\mathsf{poly}(n)$ such that $B=ds$. It follows by Lemma 4.1 that

$$\|\mathsf{Sim}(0)\,|\mathbf{0}\rangle - \mathsf{Enc}(y)\,|\mathbf{0}\rangle\,\|\leq\delta+\frac{2(m-p+1)}{d+1}+2\varepsilon\leq\delta+\frac{1}{q}\,.$$

Similarly, for every $y\in\Pi_Y$ we have $\|\mathsf{Sim}(1)\,|\mathbf{0}\rangle - \mathsf{Enc}(y)\,|\mathbf{0}\rangle\,\|\leq\delta+1/q$. This completes the proof.

□

One can also prove the following special case.

Theorem 4.2. *For any polynomial function $q(n)$, it holds that $1\mathsf{SL}[\lambda,\mu]\subseteq 1\mathsf{RE}[\delta+1/q(n),\mu]/\mathsf{poly}$ and $1\mathsf{QSL}[\lambda,\mu]\subseteq 1\mathsf{QRE}[\delta+1/q(n),\mu]/\mathsf{poly}$ where $\delta:=\min\{\sqrt{(\lambda\ln 2)/2},1-2^{-\lambda-2}\}$ and the advice size is a polynomial function of q.*

Proof. The proof is exactly similar to that of Theorem 4.1 for the special case where f is the identity function. Note that since the reduction is one-sided sparsely lossy, then Disguising Lemma 4.1 only guarantees Item 2 regarding $S_1=\Pi_Y$. Therefore, we obtain a one-sided randomized encoding. □

4.1 Better Analysis for Small λ

A different variant of Theorem 4.1 and 4.2 can be obtained for sufficiently small values of λ.

Theorem 4.3. *For any permutation-invariant function f, we have*

$$\mathsf{SL}_f[\lambda,\mu]\subseteq\mathsf{RE}[\sqrt{2\ln 2\cdot m\lambda},\mu]/\mathsf{poly},\quad\textit{and}\quad\mathsf{QSL}_f[\lambda,\mu]\subseteq\mathsf{QRE}[\sqrt{2\ln 2\cdot m\lambda},\mu]/\mathsf{poly}\,.$$

Moreover, it holds that

$$1\mathsf{SL}[\lambda,\mu]\subseteq 1\mathsf{RE}[\sqrt{2\ln 2\cdot\lambda},\mu]/\mathsf{poly},\quad\textit{and}\quad 1\mathsf{QSL}[\lambda,\mu]\subseteq 1\mathsf{QRE}[\sqrt{2\ln 2\cdot\lambda},\mu]/\mathsf{poly}\,.$$

This result does not depend on the disguising lemma nor the minimax theorem, but rely on a different information-theoretic argument. Its advantage is that when $\lambda(n)$ is negligible for a sparsely lossy problem, then the privacy of its upper-bounds are also negligible. This is in contrast with Theorem 4.1 and 4.2 where the privacy is $\delta + 1/q$, which remains inverse polynomial.

Proof. We prove the first inclusion, the others are similar. Let $\Pi \in \mathsf{SL}_f[\lambda, \mu]$. Take as advice $x_1^* \in \Pi_Y \cap \{0,1\}^n$ and $x_0^* \in \Pi_N \cap \{0,1\}^n$ (in the one-sided variant, only x_1^* suffices) and $p(f)$ be as in Lemma 2.7. The encoding and the simulator are as follows:

$$\mathsf{Enc}(y) := R\left(x_1^*, \cdots, x_1^*, x_0^*, \cdots, x_0^*, y\right),$$

and

$$\forall b \in \{0,1\} : \mathsf{Sim}(b) := R\left(x_1^*, \cdots, x_1^*, x_0^*, \cdots, x_0^*, x_b^*\right),$$

where there are $p(f)-1$ number of x_1^* and $m - p(f)$ number of x_0^*. There correctness follows immediately by construction.

We prove privacy. Fix an input y and let $b := \chi_\Pi(y)$ (for one-sided, b is always 1). Consider the following random variable $X := (X_1, \cdots, X_m)$. For all $i < m$, set X_i to be equal to the i-th entry of Enc as above and set X_m to be y with probability $1/2$ and x_b^* otherwise. Note that its distribution is splitting, pair-wise independent and 2-uniform as per Definition 13. Since $\Pi \in \mathsf{SL}_f[\lambda, \mu]$, we have that $I(X; R(X)) \leq m\lambda$. Note also that the distribution of $R(X)$ follows $\mathsf{Enc}(y)$ or $\mathsf{Sim}(b)$, each with probability $1/2$. Let M be a random variable that outputs $\mathsf{Enc}(y)$ with probability $1/2$ and $\mathsf{Sim}(b)$ otherwise. We have

$$\begin{aligned}
I(X; R(X)) &= I(X_m; R(X)) \quad \text{(since } X_1, \cdots, X_m \text{ are fixed)} \\
&= \mathbb{E}_{X_m}\left[D_{\mathrm{KL}}\left(p_{R(X)|X_m} \| p_{R(X)}\right)\right] \\
&= \frac{1}{2} D_{\mathrm{KL}}(p_{\mathsf{Enc}(y)} \| p_M) + \frac{1}{2} D_{\mathrm{KL}}(p_{\mathsf{Sim}(b)} \| p_M) \quad \text{(Jensen-Shannon Divergence)} \\
&\geq \frac{\|\mathsf{Enc}(y) - \mathsf{Sim}(b)\|^2}{2 \ln 2} \quad \text{(Pinsker's inequality)},
\end{aligned}$$

where we used the fact that $\|\mathsf{Enc}(y) - M\| = \|\mathsf{Sim}(b) - M\| = \|\mathsf{Enc}(y) - \mathsf{Sim}(b)\|/2$. The statement then follows by noting that $I(X; R(X)) \leq m\lambda$. The proof in the quantum setting is exactly similar with the difference being the use of quantum relative entropy instead of KL divergence. □

5 Randomized Encodings Are Sparsely Lossy

In this section we analyze the sparse lossiness of randomized encodings. Main theorem of this section is the following:

Theorem 5.1. *The following statement holds:*

$$1\mathsf{RE}[1 - 1/\mathsf{poly}(n), 1/2 - 1/\mathsf{poly}(n)] \subseteq 1\mathsf{SL}[\mathsf{negl}(n), \mathsf{negl}(n)].$$

We use the following lemma. For avoiding confusion, we differentiate a random variable X from its distribution (probability mass function) p_X, in this section. Second item can be found in [26, 58].

Lemma 5.1. *It hold that:*

1. *For any random variables X and Y, we have $I(X, Y) = \mathbb{E}_Y[D_{KL}(p_{X|Y} \| p_X)]$.*
2. *For any two distributions $p \ll q$, i.e. p is absolutely continuous with respect to q, then $D_{KL}(p\|q) \leq (\frac{\log e}{\alpha_q}) \cdot \|p - q\|^2$ where $\alpha_q = \min_x\{q(x) > 0\}$.*

Proof of Theorem 5.1. Let $\Pi \in \mathsf{1RE}[d = 1 - 1/\mathsf{poly}(n), \mu = 1/2 - 1/\mathsf{poly}(n)]$ with the one-sided randomized encoding Enc and the simulator Sim for the YES instances. Let $k(n) \geq (1 - d(n))^{-1}$ and $\ell(n) \geq (1/2 - \mu(n))^{-1}$ be polynomials. Let

$$(P, T) := \mathsf{XOR}(\mathsf{Enc}, \mathsf{Sim}, 1^k)\,,$$

where XOR is as in Lemma 2.5. Then define $(R, S) := (P^{\otimes \ell}, T^{\otimes \ell})$. The proof consists of showing that R is a reduction for Π with negligible error as per Definition 12 which is also one-sided $\mathsf{negl}(n)$-sparsely lossy as per Definition 13.

We first prove the lossiness. Let $B := \mathsf{poly}(n)$ and X be a random variable whose distribution is B-uniform over n-bit strings and splits over (Π_Y, Π_N) as per Definition 10. We bound the quantity $I(X; R(X))$ from above.

For any random variable Z let p_Z denote its distribution (probability mass functions). We first rewrite the mutual information in terms of Kullback-Leibler divergence as in Lemma 5.1

$$I(X; R(X)) = \mathbb{E}_{R(X)}[D_{\mathrm{KL}}\left(p_{X|R(X)} \parallel p_X\right)] \tag{6}$$

By Lemma 5.1, the KL divergence of two distributions decreases as their trace distance does, in particular

$$D_{\mathrm{KL}}\left(p_{X|R(X)} \parallel p_X\right) \leq \left(\frac{\log e}{\alpha_{p_X}}\right) \cdot \|p_{X|R(X)} - p_X\|^2$$

where $\alpha_{p_X} = \min_x\{p_X(x) > 0\}$. Substituting above in Equation (6), we obtain:

$$I(X; R(X)) \leq \left(\frac{\log e}{\alpha_{p_X}}\right) \cdot \mathbb{E}_{R(X)}\left[\|p_{X|R(X)} - p_X\|^2\right]\,. \tag{7}$$

We split the bound on the right-hand side of the Inequality (7) into two terms.

Bounding $\mathsf{term}_1 = -\log(\alpha_X)$: Since X_i's are pairwise independent B-uniform distributions, we have $\alpha_X = 1/B$.
Bounding $\mathsf{term}_2 = \mathbb{E}_{R(X)}[\|p_{X|R(X)} - p_X\|^2]$: Firstly, we have

$$\begin{aligned}
\|p_{X|R(X)} - p_X\| &= \frac{1}{2}\sum_{x,y} \Pr(R(X) = y) \cdot |\Pr\left(X = x | R(X) = y\right) - \Pr(X = x)| \\
&= \frac{1}{2}\sum_{x,y} \Pr(R(X) = y) \cdot \left|\frac{\Pr\left(X = x \wedge R(X) = y\right)}{\Pr(R(X) = y)} - \Pr(X = x)\right| \\
&= \frac{1}{2}\sum_{x,y} |\Pr\left(X = x \wedge R(X) = y\right) - \Pr(X = x) \cdot \Pr(R(X) = y)| \\
&= \|p_{(X,R(X))} - p_X \otimes p_{R(X)}\|\,.
\end{aligned}$$

With a similar argument, we obtain that $\|p_{X|R(X)} - p_X\| = \|p_{R(X)|X} - p_{R(X)}\|$. It follows that

$$\begin{aligned}
\|p_{X|R(X)} - p_X\| &= \|p_{R(X)|X} - p_{R(X)}\| \\
&= \sum_x \Pr(X = x) \cdot \|p_{R(x)} - p_{R(X)}\| \\
&\leq \max_x \|p_{R(x)} - p_{R(X)}\|\,.
\end{aligned}$$

Recall that (R, S) is obtained by applying the XOR and then direct product to Enc and Sim. By the properties of the randomized encoding, for any $x \in \Pi_Y \cap \{0,1\}^n$, it holds that

$$\|p_{R(x)} - p_{S(x)}\| \leq \ell \cdot d^k \leq \mathsf{negl}(n)\,,$$

where we used Lemma 2.5 and 2.4. Therefore $\|p_{R(X)} - p_{S(x)}\| \leq \mathsf{negl}(n)$ for any distribution of X over $\Pi_Y \cap \{0,1\}^n$. We conclude that for any $x \in \Pi_Y \cap \{0,1\}^n$ and any distribution of X over $\Pi_Y \cap \{0,1\}^n$, it holds that $\|p_{R(x)} - p_{R(X)}\|_1 \leq \mathsf{negl}(n)$. Therefore, we have

$$\begin{aligned}
\mathbb{E}_{R(X)}\left[\|p_{X|R(X)} - p_X\|^2\right] &\leq \mathbb{E}_{R(X)}\left[\|p_{X|R(X)} - p_X\|\right]^2 \quad \text{(Jensen's inequality)} \\
&\leq \mathbb{E}_{R(X)}\left[\max_x \|p_{R(x)}, p_{R(X)}\|\right]^2 \\
&\leq \mathsf{negl}(n)\,.
\end{aligned}$$

Combining upper bounds on term_1 and term_2, we obtain $I(X; R(X)) \leq 4\log(e) \cdot \mathsf{negl}(n) \cdot B$ for any splitting B-uniform distribution X of YES instances. For every polynomial B, it follows that $4\log(e) \cdot \mathsf{negl}(n) \cdot B = \mathsf{negl}(n)$ which proves the one-sided sparse lossiness.

It remains to note that R is a reduction as in Definition 12. This is implied by the correctness of the randomized encoding as follows. First, note that by correctness, for every $x \in \Pi_N$ and $y \in \Pi_Y$, we have that $\|p_{\mathsf{Enc}(x)} - p_{\mathsf{Enc}(y)}\| \geq 1 - 2\mu(n)$. Therefore, for every $x \in \Pi_N$ and $y \in \Pi_Y$, Lemma 2.6 implies that $\|p_{P(x)} - p_{P(y)}\| \geq 1/2 - \mu(n)$. Now, by using Lemma 2.4, we have

$$\begin{aligned}
\|p_{R(x)} - p_{R(y)}\| &= \|p_{P(x)}^{\otimes \ell} - p_{P(y)}^{\otimes \ell}\| \\
&\geq 1 - 2\exp\left(-\frac{\ell \cdot \|p_{P(x)} - p_{P(y)}\|^2}{2}\right) \\
&\geq 1 - \mathsf{negl}(n)\,.
\end{aligned}$$

Therefore, for every $x \in \Pi_N$ and $y \in \Pi_Y$, there exists a distinguisher $\mathcal{M}_{x,y}$ that tells $R(x)$ apart from $R(y)$ with probability $1 - \mathsf{negl}(n)$. By Lemma A.1, there exists an $\mathcal{M}$ that for every $x \in \Pi_N$ and $y \in \Pi_Y$ distinguishes $R(x)$ from $R(y)$ with probability $1 - \mathsf{negl}(n)$ (note that the quantifiers are swapped). This $\mathcal{M}$ satisfies the requirement in Lemma 12, which completes the proof.

□

One might hope to obtain a similar statement for two-sided randomized encodings and SL. However, the above technique cannot be used. The barrier comes from the fact that the reduction is XOR of the encoding $\mathsf{Enc}(\cdot)$ and the simulator $\mathsf{Sim}(\cdot)$ for the YES instances. In the two-sided setting, in order to keep the reduction lossy, we would have to use the correct simulator for the instance, i.e., $\mathsf{Sim}(\chi_\Pi(\cdot))$. Then the reduction could possibly go beyond the polynomial time. On the positive side, we can reuse the above argument when the distance d and error μ are both negligible. In this setting, there is no need to amplify these quantities via XORing and repeating. The reduction is in fact Enc itself.

Theorem 5.2. *It holds that* $\mathsf{RE}[\mathsf{negl}(n), \mu] \subseteq \mathsf{SL}[\mathsf{negl}(n), \mu]$.

We obtain a quantum variant of Theorem 5.1. The construction of the randomized encoding is the same but the proof is slightly different. This is because lack of a reverse-Pinsker-like inequality for the quantum states. Instead, we use Alicki–Fannes–Winter Inequality 5.3 which gives an upper bound for conditional Shannon entropy in terms of the trace distance.

Theorem 5.3. *The following statement holds:*

$$1\mathsf{QRE}[1 - 1/\mathsf{poly}(n), 1/2 - 1/\mathsf{poly}(n)] \subseteq 1\mathsf{QSL}[\mathsf{negl}(n), \mathsf{negl}(n)]\,.$$

We first recall the required preliminaries. We define the quantum conditional entropy of a two-system state ρ_{AB} as follows

$$S(A|B) := S(\rho_{AB}) - S(\rho_B)\,,$$

The quantum mutual information in terms of conditional quantum entropy is

$$I(A;B)_\rho = S(\rho_A) - S(A|B)_\rho = S(\rho_B) - S(B|A)_\rho \,.$$

Lemma 5.2. *Let ρ_{AB} be a quantum state in two subsystems A and B, with marginal states ρ_A and ρ_B. The following properties hold.*

1. *The Von Neumann entropy is additive for tensor product states: $S(\rho_A \otimes \rho_B) = S(\rho_A) + S(\rho_B)$.*
2. *Conditioning does not increase entropy: $S(\rho_A) \geq |S(A|B)_\rho|$.*
3. *The quantum entropy of a system is bounded by the dimension: $S(\rho_A) \leq \dim(H_A)$.*

Lemma 5.3 (Alicki–Fannes–Winter Inequality [63]). *Given two bipartite quantum states $\rho_{AB}, \omega_{AB} \in \mathcal{D}(H_A \otimes H_B)$, then*

$$|S(A|B)_\rho - S(A|B)_\omega| \leq 2\|\rho - \sigma\| \log \dim(H_A) + h(\|\rho - \sigma\|) \,,$$

where $h(p) := -p \log p - (1-p)\log(1-p)$ is the binary entropy function.

Proof of Theorem 5.3. Let $\Pi \in 1\mathsf{QRE}[d = 1 - 1/\mathsf{poly}(n), \mu = 1/2 - 1/\mathsf{poly}(n)]$ with the one-sided quantum randomized encoding Enc and the simulator Sim for the YES instances. Let $k(n) \geq (1 - d(n))^{-1}$ and $\ell(n) \geq (1/2 - \mu(n))^{-1}$ be polynomials. Let

$$(P, T) := \mathsf{XOR}(\mathsf{Enc}, \mathsf{Sim}, 1^k) \,,$$

where XOR is as in Lemma 2.5. Then define $(R, S) := (P^{\otimes \ell}, T^{\otimes \ell})$. We show that R is a quantum reduction for Π that is one-sided $\mathsf{negl}(n)$-sparsely lossy as per Definition 13 and has negligible error as per Definition 12.

Let us denote by $\rho_{X,R(X)}$ (or simply by ρ) the joint system of the classical-quantum state after the reduction R is applied to distribution X, see Equation (5). We denote the subsystems of ρ by ρ_X and $\rho_{R(X)}$. Note that since ρ is a classical-quantum system, so is $\rho_X \otimes \rho_{R(X)}$. Let also ω be the product state $\omega_{X,R(X)} := \rho_X \otimes \rho_{R(X)}$, since the mutual information between subsystems of product states are zero, we have

$$\begin{aligned} I(X;R(X))_\rho &= |I(X;R(X))_\rho - I(X;R(X))_\omega| \\ &= |S(\rho_X) - S(X|R(X))_\rho - S(\omega_X) + S(X|R(X))_\omega| \\ &= |S(X|R(X))_\rho - S(X|R(X))_\omega| \\ &\leq 2\|\rho - \omega\| \log \dim(\mathrm{Im}(R(X))) + h(\|\rho - \omega\|) \,, \end{aligned}$$

where in the last inequality we used Lemma 5.3. We have

$$\|\rho - \omega\| = \|\rho_{X,R(X)} - \rho_X \otimes \rho_{R(X)}\| = \sum_x \Pr(X = x)\|\rho_{R(x)} - \rho_{R(X)}\| \,.$$

By following the exact same argument as in the proof of Theorem 5.1, we obtain that $\|\rho_{R(x)} - \rho_{R(X)}\| \leq 2\ell d^k$ for every $x \in \Pi_Y$. It then follows that

$$\begin{aligned} I(X;R(X))_\rho &\leq 2\|\rho - \omega\| \log \dim(\mathrm{Im}(R(X))) + h(\|\rho - \omega\|) \\ &\leq 4\ell \cdot d^k \log \dim(\mathrm{Im}(R(X))) + h(2\ell \cdot d^k) \\ &\leq \mathsf{negl}(n) \,, \end{aligned}$$

where we used the fact that $\dim(\mathrm{Im}(R(X))) = \mathsf{poly}(n)$ and h is an increasing function on inputs $< 1/2$. This completes the privacy part.

The correctness is similar to the proof of Theorem 5.1.

□

6 Randomized Encodings and One-Wayness

In this section, we investigate the connections between one-sided randomized encodings, OWFs and OWSGs. The first is about the existence of OWFs.

Theorem 6.1. *If* $\mathsf{1RE}[1-1/\mathsf{poly}(n),0] \not\subseteq \mathsf{P/poly}$ *then OWFs exist.*

Proof. We prove OWFs exist if $\mathsf{1RE}[1/2-1/\mathsf{poly}(n),0] \not\subseteq \mathsf{BPP}$. Then by Theorem 4.2 and 5.1, we have that $\mathsf{1RE}[1/2-1/\mathsf{poly}(n),0] = \mathsf{1RE}[1-1/\mathsf{poly}(n),0]$ (up to some advice). Moreover, it holds that $\mathsf{BPP/poly} \subseteq \mathsf{P/poly}$ via standard repeat-then-majority derandomization. These two then yield the theorem.

Let Π be a problem in $\mathsf{1RE}[d := 1/2 - 1/n^c, 0] \setminus \mathsf{BPP}$ for some constant c. Let Enc and Sim be its one-sided encoding and simulator. Let $T(n) = \mathsf{poly}(n)$ be an upper bound on the runtime of Enc and Sim. Without loss of generality, we can assume that the size of the randomness of these circuits is $\kappa(n) \leq T(n)$ bits in both cases.[1] Let x be an arbitrary YES instance of Π and $(C_0(\cdot), C_1(\cdot))$ be $(\mathsf{Sim}(\cdot), \mathsf{Enc}(x;\cdot))$ that take $r \in \{0,1\}^\kappa$ as input. We show that $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\} \to \{0,1\}^*$, defined as $\mathsf{F}(r,b) := C_b(r)$ is a weak OWF.

Assume that there exists an adversary $\mathcal{A}$ that breaks the scheme above with probability more than $1 - 1/(2n^c)$. We use $\mathcal{A}$ to construct an algorithm for Π. Consider the following oracle algorithm $\mathcal{B}^{\mathcal{A}}$:

- $\mathcal{B}^{\mathcal{A}}(y)$: Let $\widehat{C}_0(\cdot) := \mathsf{Sim}(\cdot)$ and $\widehat{C}_1(\cdot) := \mathsf{Enc}(y;\cdot)$. Sample a uniform $r \in \{0,1\}^\kappa$ and a uniform $b \in \{0,1\}$, and let $z := \widehat{C}_b(r)$. Run the adversary $(r', b') \leftarrow \mathcal{A}(z)$, and compute $z' = C_{b'}(r)$. If $z = z'$ output 1, otherwise output 0.

$\mathcal{B}$ mimics the behavior of F but it replaces $(C_0, C_1) := (\mathsf{Sim}, \mathsf{Enc}(x))$ with $(\widehat{C}_0, \widehat{C}_1) = (\mathsf{Sim}, \mathsf{Enc}(y))$ in the state generation algorithm. We compute the advantage of $\mathcal{B}$ in distinguishing between YES and NO instances of Π by analyzing the probability $\Pr\big[\mathcal{B}^{\mathcal{A}}(y) = 1\big]$.

Case $y \in \Pi_N$: We show that for every $y \in \Pi_N$, we have:

$$\Pr\big[\mathcal{B}^{\mathcal{A}}(y) = 1\big] \leq \frac{1}{2}\,.$$

If $b = 0$, then $\mathcal{A}$ succeeds to find an inverse with probability at most 1. Therefore, $\mathcal{B}$ outputs 1 with probability at most 1. On the other hand, if $b = 1$ then the adversary is fed with $\mathsf{Enc}(y)$. By the correctness property, we know that the distance between $C_1 = \mathsf{Enc}(x)$ and $\widehat{C}_1 = \mathsf{Enc}(y)$ is 1. This means that the images of C_1 and $\widehat{C}_1$ are distinct. Therefore, there exists no r' such that $\widehat{C}_1(r) = C_1(r')$. So the adversary never succeeds and $\mathcal{B}$ never outputs 1 in this case. This yields the above upper bound.

Case $y \in \Pi_Y$: By the privacy property, the distance of the outcomes of $C_1 = \mathsf{Enc}(x)$ and $\widehat{C}_1 = \mathsf{Enc}(y)$ is at most $2d$ since both of them are YES instances. Moreover, $\widehat{C}_0$ is exactly the same as C_0. Therefore, if the bit b, chosen by $\mathcal{B}$ is equal to 0, then $\mathcal{A}$ succeeds with probability at least $1 - 1/(2n^c)$, and if $b = 1$, it succeeds with probability at least $1 - 1/(2n^c) - 2d$. In total, we obtain

$$\Pr\big[\mathcal{B}^{\mathcal{A}}(y) = 1\big] \geq \frac{1}{2}(1 - \frac{1}{2n^c}) + \frac{1}{2}(1 - \frac{1}{2n^c} - 2d) = 1 - \frac{1}{2n^c} - d \geq \frac{1}{2} + \frac{1}{2n^c}\,.$$

Reduction: The quantity of $\Pr\big[\mathcal{B}^{\mathcal{A}}(y) = 1\big]$ is different for YES and NO instances of y with an inverse polynomial gap $\mathsf{gap} := 1/(2n^c)$. To decide Π it suffices to run $\mathcal{B}$ for $\mathsf{poly}(1/\mathsf{gap})$ many times, and approximate the quantity above within error less than $\mathsf{gap}/2$. If this value is more than $1/2 + \mathsf{gap}/2$, then y must be a YES instance, otherwise it is a NO instance. We

[1] It is not difficult to show that if $\kappa = O(\log n)$ then $\Pi \in \mathsf{BPP}$.

precise the idea in the following. Let $k \in \mathbb{N}$. On instance $y \in \{0,1\}^n$, independently run $\mathcal{B}(y)$ for k times and check whether its output correctly guesses the internal random bit b or not. Let $b_1, \ldots, b_k$ be k corresponding bits representing whether the checks were passed or not. By definition, the average of $b_1, \cdots, b_k$ must be an estimation of the success probability of $\mathcal{B}$ in guessing the internal random bit b. In other words, it approximates the quantity $1/2+\mathsf{gap}/2$. Then return as follows:

$$\begin{cases} 0 & \text{if } \left|\frac{1}{k}\sum_i b_i - \frac{1}{2} + \frac{1}{4n^c}\right| \geq \frac{\mathsf{gap}}{2}, \\ 1 & \text{otherwise}\,. \end{cases}$$

To analyze the success probability of the algorithm above in deciding y, we proceed as follows. Let $\mathcal{C}$ be the algorithm. We have

$$\begin{aligned}
&\Pr(\mathcal{C}(y) = 0 | y \in \Pi_Y) \\
&= \Pr\left(\left|\frac{1}{k}\sum_i b_i - \frac{1}{2} + \frac{1}{4n^c}\right| \geq \frac{\mathsf{gap}}{2} \;\middle|\; b_1, \ldots, b_k, y \in \Pi_Y\right) \\
&\geq \Pr\left(\frac{1}{k}\sum_i b_i \geq \frac{1}{2} - \frac{1}{4n^c}\frac{\mathsf{gap}}{2} \;\middle|\; b_1, \ldots, b_k, y \in \Pi_Y\right) \\
&\geq \Pr\left(\frac{1}{k}\left(\sum_i b_i - \mathbb{E}(b_i)\right) \geq -\frac{\mathsf{gap}}{2} \;\middle|\; b_1, \ldots, b_k, y \in \Pi_Y\right) \\
&\geq 1 - \exp\left(-\frac{k \cdot \mathsf{gap}^2}{2}\right)\,,
\end{aligned}$$

where we used $\mathbb{E}(b_i) - \mathsf{gap}/2 \geq 1/2 - 1/(4n^c) - \mathsf{gap}/2$ for $y \in \Pi_Y$ as shown earlier for the second inequality, and used Hoeffding's lemma in the last inequality. On the other hand, we have

$$\begin{aligned}
&\Pr(\mathcal{C}(y) = 1 | y \in \Pi_N) \\
&= \Pr\left(\left|\frac{1}{k}\sum_i b_i - \frac{1}{2} + \frac{1}{4n^c}\right| < \frac{\mathsf{gap}}{2} \;\middle|\; b_1, \ldots, b_k, y \in \Pi_N\right) \\
&= \Pr\left(\left|\frac{1}{k}\sum_i b_i - \frac{1}{k}\sum_i \mathbb{E}(b_i)\right| < \frac{\mathsf{gap}}{2} \;\middle|\; b_1, \ldots, b_k, y \in \Pi_N\right) \\
&\geq 1 - \exp\left(-\frac{k \cdot \mathsf{gap}^2}{2}\right)\,.
\end{aligned}$$

For $k := 4/\mathsf{gap}^2$, any sufficiently large $n \in \mathbb{N}$, and any $y \in (\Pi_Y \cup \Pi_N) \cap \{0,1\}^n$, it holds that

$$\Pr(\mathcal{C}(y) = \chi_\Pi(y)) \geq 1 - \exp\left(-\frac{k \cdot \mathsf{gap}^2}{2}\right) \geq \frac{2}{3}\,,$$

Therefore, $\mathcal{C}$ solves Π which is a contradiction. Therefore, no polynomial-time adversary $\mathcal{A}$ has advantage larger than $1 - 1/(2n^c)$ in inverting F.

□

As mentioned in the introduction, building OWFs from worst-case hardness of SZK is of great interest. The above theorem proposes a possible approach for doing so. The missing puzzle is a procedure that allows to completely remove the error from one-side randomized encodings even it means to increase the privacy. We prove the following theorem.

Theorem 6.2. *If* $1\mathsf{RE}[\mathsf{negl}(n), \mathsf{negl}(n)] \subseteq 1\mathsf{RE}\left[1 - 1/\mathsf{poly}(n), 0\right]/\mathsf{poly}$ *and* SZK *is worst-case hard, then OWFs exist.*

Proof. First note that if SZK is worst-case hard then NISZK is worst-case hard by Lemma 2.10. We also have $\mathsf{NISZK} \subseteq 1\mathsf{RE}[\mathsf{negl}(n), \mathsf{negl}(n)]/\mathsf{poly}$ by Lemma 2.11. The rest follows from Theorem 6.1 and the assumption. □

Next theorem is the adaptation of Theorem 6.1 for OWSGs. The construction is similar. We let the generation algorithm of the OWSG to be the quantum mapping $\mathsf{StateGen}$ that maps (b, r) to $\mathsf{Sim}(r)$ if $b = 0$ and $\mathsf{Enc}(x; r)$ if $b = 1$, where x is an arbitrary YES instance. The verification algorithm takes (r, b) and a quantum state ρ as inputs and accepts if $\|\mathsf{StateGen}(r, b) - \rho\| \leq 1/3$.

Theorem 6.3. *If* $1\mathsf{QRE}_{\mathsf{pure}}[1 - 1/\mathsf{poly}(n), 0] \not\subseteq \mathsf{BQP}/\mathsf{poly}$ *then OWSGs exist.*

Proof. It suffices to prove the theorem for $1\mathsf{QRE}_{\mathsf{pure}}[1/2 - 1/\mathsf{poly}(n), 0] \not\subseteq \mathsf{BQP}$ since by Theorem 4.2 and 5.1, we have that $1\mathsf{QRE}_{\mathsf{pure}}[1/2 - 1/\mathsf{poly}(n), 0] = 1\mathsf{QRE}_{\mathsf{pure}}[1 - 1/\mathsf{poly}(n), 0]$ (up to some polynomial-size classical advice).[1] Let $\Pi \in 1\mathsf{QRE}_{\mathsf{pure}}[d = 1/2 - 1/n^c, 0] \setminus \mathsf{BQP}$ for some constant c with one-sided encoding Enc and simulator Sim for YES instances of Π. Let x be an arbitrary YES instance of Π. Let (C_0, C_1) be $(\mathsf{Sim}, \mathsf{Enc}(x))$. Note that the two circuits are mixed; a classical randomness is used. Let κ be the size of the randomness of these circuits.[2] We show that G, defined as follows is a weak OWSG:

- $\mathsf{StateGen}(r, b)$: output $C_b |r, \mathbf{0}\rangle$.
- $\mathsf{Ver}((r, b), \rho)$: If $\|C_b |r, \mathbf{0}\rangle - \rho\|_1 \leq 1/3$ output 1, otherwise output 0.

This suffices to build OWSG according to [50].

Assume that there exists an adversary $\mathcal{A}$ that breaks the scheme above with probability more than $1 - 1/(2n^c)$ and let $\mathcal{B}^{\mathcal{A}}$ be the following oracle algorithm:

- $\mathcal{B}^{\mathcal{A}}(y)$:
  - State Generation Phase: Let $\widehat{C}_0(\cdot) := \mathsf{Sim}(\cdot)$ and $\widehat{C}_1(\cdot) := \mathsf{Enc}(y; \cdot)$. Sample a uniform $r \in \{0,1\}^\kappa$ and a uniform $b \in \{0,1\}$, and let $\rho := \widehat{C}_b |r, \mathbf{0}\rangle$.
  - Test Phase: Run the adversary $(r', b') \leftarrow \mathcal{A}(\rho^{\otimes n})$, and compute $\rho' = C_{b'} |r', \mathbf{0}\rangle$. If $\|\rho - \rho'\| \leq 1/3$ output 1, otherwise output 0.

Similar as in Theorem 6.1, one can show that $\Pr\big(\mathcal{B}^{\mathcal{A}}(y) = 1\big) \leq 1/2$ if $y \in \Pi_N$, and $\Pr\big(\mathcal{B}^{\mathcal{A}}(y) = 1\big) \geq 1/2 + 1/(2n^c)$ if $y \in \Pi_Y$. To decide Π it suffices to run $\mathcal{B}$ for $\mathsf{poly}(n^c)$ many times, and approximate the quantity above within error less than $1/(4n^c)$. If this value is more than $1/2 + 1/(4n^c)$, then y must be a YES instance, otherwise it is a NO instance. Therefore, we finally obtain a algorithm that solves Π. Note that $\mathcal{B}$ tests whether $\|\widehat{C}_b |r, \mathbf{0}\rangle - \widehat{C}_{b'} |r', \mathbf{0}\rangle\|_1$ is smaller than $1/3$. Since the encoding Enc is pure and r, r' are fixed, these states are pure, therefore $\mathcal{B}$ can perform a SWAP test for $\mathsf{poly}(n)$ number of times on them to approximate their distance. □

7 Obfuscation

We start this section by defining the obfuscation functionality as below.

Definition 16 (Statistical Obfuscation (sO), [36]). *Let N be a positive integer, and $\alpha, \varepsilon : \mathbb{N} \to [0, 1)$ be functions. An α-statistical obfuscation with error ε is an algorithm sO that, upon receiving a security parameter 1^N and a circuit C of size $\mathsf{poly}(N)$ as inputs, runs in time $\mathsf{poly}(N)$ and outputs a circuit with the following specifications:*

[1] The statements of Theorem 4.2 and 5.1 does not explicitly address $1\mathsf{QRE}_{\mathsf{pure}}$, however the transformations that are involved in the proofs respect the purity of the encodings and reductions.

[2] It is not difficult to show that if $\kappa = O(\log n)$ then $\Pi \in \mathsf{BQP}$.

1. ***Correctness.*** *For any circuit C over inputs of size N, it holds that*

$$\Pr\left[\exists x \in \{0,1\}^N \ : \ \mathsf{sO}(1^N, C)(x) \neq C(x)\right] \leq \varepsilon\,,$$

where the probability is taken over the randomness of sO.

2. ***Statistical Distance.*** *For all circuits C_1 and C_2 that are functionally equivalent over inputs of size N such that $|C_1| = |C_2|$, it holds that*

$$\Delta(\mathsf{sO}(1^N, C_1), \mathsf{sO}(1^N, C_2)) \leq \alpha(N)\,. \tag{8}$$

Remark 6 (Indistinguishability obfuscation). An indistinguishability obfuscation (iO) is defined similarly where the error is $\mathsf{negl}(N)$ and Equation (8) is replaced by computational indistinguishability: No polynomial-time adversary can distinguish between the two obfuscations with probability more than $\mathsf{negl}(N)$. We say that iO is weak if the computational distinguishing probability is bounded by $1 - 1/\mathsf{poly}(N)$.

The main theorem of this section the following one.

Theorem 7.1. *Assuming that Polynomial Hierarchy does not collapse to its third level, $(1 - 1/\mathsf{poly}(N))$-statistical obfuscation with error $1/2 - 1/\mathsf{poly}(N)$ does not exist.*

Assuming that refuting 3SAT *is exponentially hard for non-uniform algorithms, $(1 - 1/\mathsf{subexp}(N))$-statistical obfuscation with error $1/2 - 1/\mathsf{subexp}(N)$ does not exist.*

The assumption that refuting 3SAT is exponentially hard is also known as non-deterministic exponential time hypothesis. It is is introduced by [21], and more formally states that 3SAT does not have a non-uniform reduction to coNP that runs in time subexponential in N, where N is the number of variables in the 3SAT instance. It has been further studied by [21, 23].

Remark 7. We note that this impossibility result should not be confused with the positive result of Brakerski, Brzuska, and Fleischhacker [17]. They construct an efficient obfuscator when $2\varepsilon + \alpha > 1$ [17, Appendix A], but their correctness definition of obfuscation is quite different. They require that the obfuscator correctly evaluates the original circuit only over a uniformly random input, which is a relatively weaker guarantee.

The following is an immediate application of the main theorem.

Corollary 7.1. *If Polynomial Hierarchy does not collapse to its third level, then weak iO implies OWFs. If refuting* 3SAT *is exponentially hard for non-uniform algorithms, then weak iO with subexponential time implies OWFs.*

Proof. Recall that iO is an obfuscation scheme with negligible error $\varepsilon(N) = \mathsf{negl}(N)$ such that the obfuscation of any two functionally equivalent circuits are computationally indistinguishable. On the other hand, in Theorem 7.1 implies that there are two family of circuits whose obfuscations are statistically far assuming that Polynomial Hierarchy does not collapse. Consider these obfuscations. They are computationally close but statistically far with a gap that is inverse polynomially large. This is known as EFI pairs which imply OWFs [12, 33, 51]. The proof for the second part is similar and can be obtained by changing the parameter N. □

In the rest of this section, we prove the main result. We need the following lemma. We are interested in the unique satisfiability problem (UNIQUESAT) which is similar to 3SAT with the extra promise that the input 3CNF formula has at most one satisfiable assignment. We also consider the regime of parameters where the runtime of the encoding and simulator can be superpolynomial since it has interesting consequences.

Lemma 7.1. *Assume that there exists an α-statistical obfuscation sO with error ε such that $\alpha = 1 - 1/\mathsf{poly}(N)$ and $\varepsilon = 1/2 - 1/\mathsf{poly}(N)$. Then there exists a one-sided randomized encoding for* UniqueSat *with privacy $1 - 1/\mathsf{poly}(n)$, error $1/2 - 1/\mathsf{poly}(n)$, and the same runtime as sO, where n is the size of the instance of* UniqueSat.

Proof. Let C be a YES instance of UniqueSat, i.e., it evaluates to 1 on exactly one input, then a random shift of C, namely the circuit $C_z(x) = C(x \oplus z)$, where $z \xleftarrow{\$} \{0,1\}^N$, has a truth-table whose distribution is identical to that of a random point function. We let the encoding to be this random shift followed by applying sO. The one-sided simulator is $\mathsf{sO}(\mathbb{1}_{x=z})$, where $\mathbb{1}_{x=z}$ is the point function on z where z is sampled uniformly at random from $\{0,1\}^N$.

We analyze the distance d, error μ of the encoding. If C is satisfiable, then the truth table of the point function $\mathbb{1}_{x=z}$, where $z \xleftarrow{\$} \{0,1\}^N$, is statistically equivalent to that of C_z. Therefore, by the statistical indistinguishability of sO, if C is satisfiable, then $\Delta(C_i, \mathsf{sO}(1^N, \mathbb{1}_{x=z})) \leq \alpha(N)$. Note that the size of the instance $n = |C|$ is $\mathsf{poly}(N)$. Without loss of generality, we can assume that $|C| = \Omega(N)$; else, we can pad C with enough garbage bits to add to its length, and throw out these extra bits in the reduction. Thus, if $\alpha(N) = 1 - 1/q(N)$ for some polynomial function q then the privacy is $\leq 1 - 1/q(n) = 1 - 1/\mathsf{poly}(n)$. The error follows immediately by the correctness of the obfuscation.

□

The above lemma shows that UniqueSat reduces to $1\mathsf{RE}[1 - 1/\mathsf{poly}(n), 1/2 - 1/\mathsf{poly}(n)]$ assuming that such an obfuscation exists. In the following, we prove that it is unlikely that UniqueSat has such a randomized encoding, and therefore, the assumption must be false.

Valiant and Vazirani [60] show that there exists a reduction from 3Sat to UniqueSat such that given an instance C (a circuit over inputs $x \in \{0,1\}^N$), runs in time $\mathsf{poly}(N)$ and has error at most $1 - 1/\mathsf{poly}(N)$. The following lemma is obtained using their reduction.

Lemma 7.2. *There is a reduction from* 3Sat *to* Maj$\circ$UniqueSat *that runs in time $\mathsf{poly}(N)$, has zero error and takes a non-uniform advice as input.*

Proof. Let C be an instance of 3Sat with size at most $q(N)$. Let R be the Valiant-Vazirani reduction and p be such that R has error at most $1 - 1/p(N)$. By repeating the reduction for a sufficiently large $k = \mathsf{poly}(p, N)$ times and taking the majority, the error reduces to $e^{-\Omega(k)}$ according to Lemma 2.1. Moreover, the union bound implies that R errs over at least one circuit C of size $q(N)$ with probability at most $2^{q(N)} \cdot e^{-\Omega(k)}$. By carefully tuning k to be a polynomial larger than q, this error quantity becomes strictly smaller than 1. Therefore, there exists at least one random string such that the reduction correctly maps all the instances C of size q. We set this random string as the non-uniform advice. Therefore, this reduction runs in $\mathsf{poly}(N)$ time and has zero error.

□

Finally, we can prove the main theorem.

Proof of Theorem 7.1. We prove the statements separately:

Proving the first statement: Assume $(1 - 1/\mathsf{poly}(N))$-statistical obfuscation with error $1/2 - 1/\mathsf{poly}(N)$ exists. Lemma 7.1 presents a one-sided randomized encoding for UniqueSat with privacy $1 - 1/\mathsf{poly}(n)$ and error $1/2 - 1/\mathsf{poly}(n)$. On the other side, we have that $\mathsf{RE}[1 - 1/\mathsf{poly}(n), 1/2 - 1/\mathsf{poly}(n)] \subseteq 1\mathsf{RE}[\mathsf{negl}(n), \mathsf{negl}(n)]$ by Theorem 4.2, 4.3 and 5.1 (or simply Key Lemma in Section 1.1). Then Lemma 2.11 implies that UniqueSat reduces to SZK in polynomial time. On the other hand, Lemma 7.2 provides a zero-error reduction from 3Sat to Maj $\circ$ UniqueSat. Combining these two reductions, we obtain a zero-error reduction from 3Sat to Maj $\circ$ SZK that runs in polynomial time. The class SZK is closed under

Boolean formulas by Lemma 2.8, therefore, the whole chain of reductions imply that 3Sat $\in$ SZK/poly. Moreover, SZK/poly $\subseteq$ coNP/poly by Lemma 2.9. Therefore, we obtain NP $\subseteq$ coNP/poly which contradicts the assumption according to Lemma 2.12.

Proving the second statement: The proof is exactly the same with the only difference being the runtime of the reduction from 3Sat to coNP/poly. In this case, the runtime is subexponential. However, this contradicts the assumption on the subexponential hardness of refuting 3Sat. □

AI Disclosure. The lower bound on the mutual information in Theorem 4.3 was calculated by Fable 5 (Anthropic) AI agent. The encoding in the theorem is itself an immediate result of Definition 13 which is generated by humans. We also used the same model to convert a hand-drawn version of Figure 1 to LaTeX code, and occasionally to polish a few parts of the text.

References

1. Shweta Agrawal, Yuval Ishai, Dakshita Khurana, and Anat Paskin-Cherniavsky. Statistical randomized encodings: A complexity theoretic view. In *Automata, Languages, and Programming*, 2015.
2. Miklós Ajtai. The shortest vector problem in ℓ_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, 1998.
3. Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on NP-hardness. STOC '06, 2006.
4. Michael Alekhnovich. More on average case vs approximation complexity. FOCS '03, 2003.
5. Benny Applebaum. Randomly encoding functions: A new cryptographic paradigm. In *Information Theoretic Security*, 2011.
6. Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0. *SIAM Journal on Computing*, 2006.
7. Benny Applebaum and Pavel Raykov. On the relationship between statistical zero-knowledge and statistical randomized encodings. In *Advances in Cryptology – CRYPTO 2016*, 2016.
8. Saikrishna Badrinarayanan, Vipul Goyal, Aayush Jain, and Amit Sahai. Verifiable functional encryption. In *Advances in Cryptology – ASIACRYPT 2016*, 2016.
9. Marshall Ball, Elette Boyle, Akshay Degwekar, Apoorvaa Deshpande, Alon Rosen, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. Cryptography from information loss. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, 2020.
10. Boaz Barak, Zvika Brakerski, Ilan Komargodski, and Pravesh K. Kothari. Limits on low-degree pseudorandom generators (or: Sum-of-squares meets program obfuscation). In *Advances in Cryptology – EUROCRYPT 2018*, 2018.
11. Boaz Barak, Samuel B. Hopkins, Aayush Jain, Pravesh Kothari, and Amit Sahai. Sum-of-squares meets program obfuscation, revisited. In *Advances in Cryptology – EUROCRYPT 2019*, 2019.
12. Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Statistical difference beyond the polarizing regime. In *Theory of Cryptography: 17th International Conference, TCC 2019, Nuremberg, Germany, December 1–5, 2019, Proceedings, Part II*, 2019.
13. Nir Bitansky, Omer Paneth, and Alon Rosen. On the cryptographic hardness of finding a nash equilibrium. FOCS '15, 2015.
14. Andrej Bogdanov and Christina Brzuska. On basing size-verifiable one-way functions on NP-Hardness. In *Theory of Cryptography Conference*, 2015.

15. Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for NP problems. *SIAM Journal on Computing*, 2006.

16. John Bostanci, Jonas Haferkamp, Dominik Hangleiter, and Alexander Poremba. Efficient Quantum Pseudorandomness from Hamiltonian Phase States. *Leibniz Int. Proc. Inf.*, 2025.

17. Zvika Brakerski, Christina Brzuska, and Nils Fleischhacker. On statistically secure obfuscation with approximate correctness. In *Advances in Cryptology – CRYPTO 2016*, 2016.

18. Zvika Brakerski and Henry Yuen. Quantum garbled circuits. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, 2022.

19. Christina Brzuska, Pooya Farshim, and Arno Mittelbach. Indistinguishability obfuscation and uces: The case of computationally unpredictable sources. In *Advances in Cryptology – CRYPTO 2014*, 2014.

20. Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Phys. Rev. Lett.*, 2001.

21. Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. ITCS '16, 2016.

22. Suvradip Chakraborty, James Hulett, and Dakshita Khurana. On weak nizks, one-way functions and amplification. In *Advances in Cryptology – CRYPTO 2025*, 2025.

23. Lijie Chen, Ron D. Rothblum, Roei Tell, and Eylon Yogev. On exponential-time hypotheses, derandomization, and circuit lower bounds. *J. ACM*, 2023.

24. Kai-Min Chung, Huijia Lin, and Rafael Pass. Constant-round concurrent zero-knowledge from indistinguishability obfuscation. In *Advances in Cryptology – CRYPTO 2015*, 2015.

25. Aloni Cohen, Justin Holmgren, Ryo Nishimaki, Vinod Vaikuntanathan, and Daniel Wichs. Watermarking cryptographic capabilities. STOC '16, 2016.

26. I. Csiszar and Z. Talata. Context tree estimation for not necessarily finite memory processes, via bic and mdl. *IEEE Transactions on Information Theory*, 2006.

27. W. Diffie and M. Hellman. New directions in cryptography. *IEEE Trans. Inf. Theory*, 1976.

28. Andrew Drucker. New limits to classical and quantum instance compression. *SIAM Journal on Computing*, 2015.

29. Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM Journal on Computing*, 1993.

30. Lance Fortnow and Rahul Santhanam. Infeasibility of instance compression and succinct pcps for np. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, 2008.

31. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *SIAM Journal on Computing*, 2016.

32. Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. Revisiting the cryptographic hardness of finding a nash equilibrium. In *Advances in Cryptology – CRYPTO 2016*, 2016.

33. Oded Goldreich. A note on computational indistinguishability. *Information Processing Letters*, 1990.

34. Oded Goldreich, Amit Sahai, and Salil Vadhan. Can statistical zero knowledge be made non-interactive? or on the relationship of szk and niszk. In *Advances in Cryptology — CRYPTO' 99*, 1999.

35. Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. In *Advances in Cryptology – EUROCRYPT 2014*, 2014.

36. Shafi Goldwasser and Guy N. Rothblum. On best-possible obfuscation. In *Theory of Cryptography*, 2007.

37. Shuichi Hirahara and Mikito Nanashima. One-way functions and zero knowledge. STOC 2024, 2024.

38. Rahul Ilango and Alex Lombardi. Cryptography meets worst-case complexity: Optimal security and more from iO and worst-case assumptions, 2025.

39. Y. Ishai and E. Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, 2000.

40. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. STOC 2021, 2021.

41. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from lpn over fp, dlin, and prgs in nc0. In *Advances in Cryptology – EUROCRYPT 2022*, 2022.

42. Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, 2024.

43. Dakshita Khurana and Kabir Tomer. Founding quantum cryptography on quantum advantage, or, towards cryptography from #p hardness. STOC '25, 2025.

44. Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yogev. One-way functions and (im)perfect obfuscation. FOCS '14, 2014.

45. William Kretschmer. Quantum Pseudorandomness and Classical Complexity. In *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, 2021.

46. William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, 2023.

47. Richard J. Lipton and Neal E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '94, 1994.

48. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 2007.

49. Tomoyuki Morimae, Yuki Shirakawa, and Takashi Yamakawa. From worst-case hardness of NP to quantum cryptography via quantum indistinguishability obfuscation. Cryptology ePrint Archive, Paper 2025/1185, 2025.

50. Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. In *19th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2024*, 2024.

51. Moni Naor and Guy N. Rothblum. Learning to impersonate. In *Proceedings of the 23rd International Conference on Machine Learning*, ICML '06, 2006.

52. R. Ostrovsky. One-way functions, hard on average problems, and statistical zero-knowledge proofs. In *[1991] Proceedings of the Sixth Annual Structure in Complexity Theory Conference*, 1991.

53. R. Ostrovsky and A. Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *The 2nd Israel Symposium on Theory and Computing Systems*, 1993.

54. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 2009.

55. R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 1978.

56. Amit Sahai and Salil P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 2003.

57. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. *SIAM Journal on Computing*, 2021.

58. Igal Sason. On reverse pinsker inequalities. *CoRR*, 2015.

59. J. v. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 1928.
60. L.G. Valiant and V.V. Vazirani. NP is as easy as detecting unique solutions. *Theoretical Computer Science*, 1986.
61. J. Watrous. Limits on the power of quantum statistical zero-knowledge. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, 2002.
62. Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious lwe sampling. In *Advances in Cryptology – EUROCRYPT 2021*, 2021.
63. Mark M Wilde. *Quantum information theory*. Cambridge university press, 2013.
64. A. Winter. Coding theorem and strong converse for quantum channels. *IEEE Transactions on Information Theory*, 1999.
65. Andrew C. Yao. Theory and application of trapdoor functions. In *23rd Annual Symposium on Foundations of Computer Science (sfcs 1982)*, 1982.
66. Chee K. Yap. Some consequences of non-uniform conditions on uniform classes. *Theoretical Computer Science*, 1983.

A Proof of Extended Disguising Lemma

We first provide necessary materials for the proof.

Two-Player games. A two-player, simultaneous-move, zero-sum game is specified by a matrix $\mathbf{M} \in \mathbb{R}^{a\times b}$. Player 1 chooses a row index $i \in [a]$ and Player 2 chooses a column index $j \in [b]$, and Player 2 receives the payoff $\mathbf{M}_{ij}$ from Player 1. The goal of Player 1 is minimizing the expected payoff, while Player 2 opts to maximize it. The row and column indices are called the pure strategies of Player 1 and Player 2, respectively. The mixed strategies are distributions or possible choices of indices. A mixed strategy is s-uniform if it is sampled uniformly from a multiset of at most s pure strategies.

Lemma A.1 ([59]). *Let $\mathcal{P}$ and $\mathcal{Q}$ be two mixed strategies for Player 1 and 2, respectively. It holds that*

$$\min_{\mathcal{P}} \max_{j} \mathbb{E}_{i\sim\mathcal{P}}[\mathbf{M}_{ij}] = \max_{\mathcal{Q}} \min_{i} \mathbb{E}_{j\sim\mathcal{Q}}[\mathbf{M}_{ij}].$$

The value of the game, which we denote by $\omega(\mathbf{M})$, is the optimal expected value guaranteed by the above lemma. The following lemma shows that each player has nearly-optimal s-uniform strategy when s is chosen to be logarithm of the number of pure strategies of the opponent.

Lemma A.2 ([47, Theorem 2]). *For any real $\varepsilon > 0$, any $\mathbf{M} \in \mathbb{R}^{a\times b}$, and any integer $s \geq \ln(b)/(2\varepsilon^2)$, it holds that*

$$\min_{\mathcal{P}\in\mathfrak{P}_s} \max_{j} \mathbb{E}_{i\sim\mathcal{P}}[\mathbf{M}_{ij}] \leq \omega(\mathbf{M}) + \varepsilon(\mathbf{M}_{max} - \mathbf{M}_{min}) \ ,$$

where $\mathfrak{P}_s$ denotes the set of all s-uniform strategies for Player 1. Similar statement holds for Player 2, namely,

$$\max_{\mathcal{Q}\in\mathfrak{Q}_s} \min_{i} \mathbb{E}_{j\sim\mathcal{Q}}[\mathbf{M}_{ij}] \geq \omega(\mathbf{M}) - \varepsilon(\mathbf{M}_{max} - \mathbf{M}_{min}) \ ,$$

where $\mathfrak{Q}_s$ denotes the set of all s-uniform strategies for Player 2.

Similar to [9, 28], we define distributional stability as follows.

Definition 17. *Let n, m, m_0, m_1 be positive integers such that $m = m_0 + m_1 + 1$. For a real $\delta \in [0,1]$, a quantum mapping $R : \{0,1\}^{mn} \to \mathsf{MS}_*$ is said to be δ-quantumly-distributionally stable (δ-QDS) with respect to two distributions $(\mathcal{D}_0, \mathcal{D}_1)$ over $\{0,1\}^n$ if the following holds:*

$$\mathop{\mathbb{E}}_{\substack{y\sim\mathcal{D}_0\\ \pi\sim\mathfrak{S}_m}}\left[\left\|R\left(\pi\left(\mathcal{D}_0^{\otimes m_0},y,\mathcal{D}_1^{\otimes m_1}\right)\right)-R\left(\pi\left(\mathcal{D}_0^{\otimes(m_0+1)},\mathcal{D}_1^{\otimes m_1}\right)\right)\right\|\right]\leq\delta\ .$$

Note that the order of the pair $(\mathcal{D}_0,\mathcal{D}_1)$ matters. Furthermore, when $m_1=0$, we simply say that the mapping is δ-QDS with respect to $\mathcal{D}_0$.

Below, we recall an adaptation of [28, Lemma 8.10].

Lemma A.3. *Assume that $R:\{0,1\}^{m\cdot n}\to\mathsf{MS}_*$ satisfies the properties in Lemma 4.1 for $m_1=0$. Then R is δ-QDS with respect to any ds-uniform distribution $\mathcal{D}_0$ supported on either S_0 or S_1.*

In the original lemma from [28], compression is used to bound the entropy of the mutual information. However, note that this can be argued directly from splitting lossiness, and that any restriction on the input distributions will give a result for the same restricted case.

The following lemma is the generalization of the above one.

Lemma A.4. *Assume that $R:\{0,1\}^{mn}\to\mathsf{MS}_*$ satisfies the properties in Lemma 4.1. Then R is δ-QDS with respect to any ds-uniform independent distributions $(\mathcal{D}_0,\mathcal{D}_1)$ each supported on either S_0 or S_1.*

Proof. The proof is similar to that of [9, Proposition B.1]. Let $\pi\in\mathfrak{S}_m$ be a fixed permutation. One can rewrite it as the composition of two partial permutations π_0 and π_1, i.e., $\pi=\pi_0\circ\pi_1$, such that π_1 only acts on the last m_1 arguments of the input. Let $\rho_\pi(y)$ be as follows

$$\rho_\pi(y):=R\left(\pi\left(\mathcal{D}_0^{\otimes m_0},y,\mathcal{D}_1^{\otimes m_1}\right)\right)\ .$$

For $y,y'\sim\mathcal{D}_0$, two independent random variables, and $\pi\sim\mathfrak{S}_m$, we want to prove that

$$\mathbb{E}_{y,\pi}\left[\left\|\rho_\pi(y)-\rho_\pi(y')\right\|\right]\leq\delta\,.$$

Note that it is enough to bound the conditional distributions since

$$\mathbb{E}_{y,\pi}\left[\left\|\rho_\pi(y)-\rho_\pi(y')\right\|\right]=\mathbb{E}_\pi\left[\mathbb{E}_{y,\pi|\pi_1}\left[\left\|\rho_\pi(y)-\rho_\pi(y')\right\|\right]\right],$$

by the law of total probability. Let $R'(x_1,x_2,\ldots,x_{m_0+1})$ be the mapping that first samples π then evaluates $R\left(\pi_1\left(x_1,x_2,\ldots,x_{m_0+1},\mathcal{D}_1^{\otimes m_1}\right)\right)$. For any fixed π_1, we show that R' is ℓ-lossy for all *ds*-uniform distributions that split over (S_0,S_1). Indeed, let $(\mathcal{X}_1,\ldots,\mathcal{X}_{m_0+1})$ be independent *ds*-uniform random variables with $\mathrm{Supp}(\mathcal{X}_i)\subseteq S_0$ or $\mathrm{Supp}(\mathcal{X}_i)\subseteq S_1$ for each $i\in[m_0+1]$, and $(\mathcal{Z}_1,\ldots,\mathcal{Z}_{m_1})\sim\mathcal{D}_1^{\otimes m_1}$, thus $\mathrm{Supp}(\mathcal{Z}_i)\subseteq\mathrm{Supp}(\mathcal{D}_1)\subseteq S_j$ for all $i\in[m_1]$ and some $j\in\{0,1\}$. By the splitting lossiness of R for any *ds*-uniform distribution, we can bound the loss of R':

$$\begin{aligned}\ell&\geq I_q(\pi_1(\mathcal{X}_1,\ldots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\ldots,\mathcal{Z}_{m_1});\\&\qquad\qquad R(\pi_1(\mathcal{X}_1,\ldots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\ldots,\mathcal{Z}_{m_1})))\\&=I_q(\mathcal{X}_1,\ldots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\ldots,\mathcal{Z}_{m_1};R(\pi_1(\mathcal{X}_1,\ldots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\ldots,\mathcal{Z}_{m_1})))\\&\geq I_q(\mathcal{X}_1,\ldots,\mathcal{X}_{m_0+1};R(\pi_1(\mathcal{X}_1,\ldots,\mathcal{X}_{m_0+1},\mathcal{Z}_1,\ldots,\mathcal{Z}_{m_1}))).\end{aligned}$$

Finally, by Lemma A.3 a splitting lossy map must also be δ-QSD, thus

$$\begin{aligned}&\mathbb{E}_{y,\pi|\pi_1}\left[\left\|\rho_\pi(y)-\rho_\pi(y')\right\|\right]\\&=\mathbb{E}_{y,\pi|\pi_1}\left[\left\|R\left(\pi\left(\mathcal{D}_0^{\otimes m_0},y,\mathcal{D}_1^{\otimes m_1}\right)\right)-R\left(\pi\left(\mathcal{D}_0^{\otimes m_0},y',\mathcal{D}_1^{\otimes m_1}\right)\right)\right\|\right]\\&=\mathbb{E}_{y,\pi_0}\left[\left\|R'\left(\mathcal{D}_0^{\otimes m_0},y\right)-R'\left(\mathcal{D}_0^{\otimes m_0},y'\right)\right\|\right]\\&\leq\delta\,.\end{aligned}$$

□

If a mapping is distributionally stable with respect to a pair of distributions, then one can "sparsify" the distributions while nearly keeping the stability.

Lemma A.5. *Let $n, m, m_0, m_1, \ell, S_0, S_1, R$ and δ be as in Lemma 4.1. Let $\mathcal{D}_0$ and $\mathcal{D}_1$ be two independent distributions with supports over S_0 and S_1, respectively. Let $\{x_i^{(0)}\}_{i\in[d+1]}$ and $\{x_i^{(1)}\}_{i\in[d+1]}$ be independent samples from $\mathcal{D}_0$ and $\mathcal{D}_1$, respectively. For each $j \in \{0,1\}$, let $y_j^* := x_{i^*}^{(j)}$ be uniformly chosen from $\{x_i^{(j)}\}_{i\in[d+1]}$ and let $\widehat{\mathcal{D}}_j$ be the uniform distribution over the multiset $\{x_i^{(j)}\}_{i\in[d+1]\setminus\{i^*\}}$. Then it holds that*

$$\mathbb{E}_{\pi\sim\mathfrak{S}_m}\left[\left\|R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\|\right] \leq \delta + \frac{2m_0+1}{d+1},$$

and

$$\mathbb{E}_{\pi\sim\mathfrak{S}_m}\left[\left\|R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_1^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, \widehat{\mathcal{D}}_1^{\otimes (m_1+1)}\right)\right)\right\|\right] \leq \delta + \frac{2m_1+1}{d+1}.$$

Proof. We prove the first statement. The other one is implied similarly. Let $\widetilde{\mathcal{D}}_0$ denote the uniform distribution over $\{x_i^{(0)}\}_{i\in[d+1]}$. For any fixed set of of multisets as above and any choice of permutation π and quantum mapping R, we have

$$\begin{aligned}
&\left\|R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\| \\
&\leq \left\|\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)} \otimes \widehat{\mathcal{D}}_1^{\otimes m_1} - \widehat{\mathcal{D}}_0^{\otimes (m_0+1)} \otimes \widehat{\mathcal{D}}_1^{\otimes m_1}\right\| \\
&\leq \left\|\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)} - \widehat{\mathcal{D}}_0^{\otimes (m_0+1)}\right\| \\
&\leq (m_0+1)\left\|\widetilde{\mathcal{D}}_0 - \widehat{\mathcal{D}}_0\right\|,
\end{aligned}$$

where we used the quantum data processing inequality for the first two upper bounds, and the property of tensor product for the last one. Since both $\widehat{\mathcal{D}}_0$ and $\widetilde{\mathcal{D}}_0$ are classical, their trace distance coincides with their statistical distance. Therefore, we have

$$\begin{aligned}
\left\|\widetilde{\mathcal{D}}_0 - \widehat{\mathcal{D}}_0\right\| &= \frac{1}{2}\sum_{x\in\{x_i^{(0)}\}_{i\in[d+1]}} \left|\Pr_{\widetilde{\mathcal{D}}_0}(x) - \Pr_{\widehat{\mathcal{D}}_0}(x)\right| \\
&= \frac{1}{2(d+1)} + \frac{1}{2}\sum_{x\in\{x_i^{(0)}\}_{i\in[d+1]\setminus\{i^*\}}} \left|\frac{1}{d+1} - \frac{1}{d}\right| \\
&= \frac{1}{d+1}.
\end{aligned}$$

Similarly, it holds that

$$\left\|R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\| \leq \frac{m_0}{d+1}.$$

From the triangle inequality, it follows that

$$\begin{aligned}
&\left\|R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\| \\
&\leq \left\|R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\| \\
&\quad + \left\|R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\| \\
&\quad\quad + \left\|R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\| \\
&< \left\|R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes m_0}, y_0^*, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widetilde{\mathcal{D}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{D}}_1^{\otimes m_1}\right)\right)\right\| \\
&\quad + \frac{2m_0+1}{d+1} \ .
\end{aligned}$$

Recall that R is ℓ-lossy with respect to all ds-uniform distributions that split over (S_0, S_1). Therefore, by Lemma A.4 it is δ-QSD with respect to all ds-uniform pair of distributions each supported on either S_0 or S_1, including $(\widetilde{\mathcal{D}}_0, \widehat{\mathcal{D}}_1)$. Finally, by taking expectation from both sides above with respect to π, and using the fact that R is δ-QSD with respect to $(\widetilde{\mathcal{D}}_0, \widehat{\mathcal{D}}_1)$, one obtains the claimed upper bound.

□

Proof of Lemma 4.1. Consider the following two-player, simultaneous-move, and zero-sum game:

- Player 1: chooses a pair of multisets $K \subseteq S_0$ and $T \subseteq S_1$, each of size d.
- Player 2: chooses an element $y \in S_0 \cup S_1$
- Payoff: if $y \in S_0$, Player 2 gains

$$\mathbb{E}_{\pi \sim \mathfrak{S}_m}\left[\left\|R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_K^{\otimes (m_0+1)}, \mathcal{U}_T^{\otimes m_1}\right)\right)\right\|\right] \ ,$$

  otherwise, Player 2 gains

$$\mathbb{E}_{\pi \sim \mathfrak{S}_m}\left[\left\|R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, \mathcal{U}_T^{\otimes (m_1+1)}\right)\right)\right\|\right] \ .$$

Consider a ds-uniform strategy for Player 2, i.e. a distribution $\mathcal{Y}$ of y that is uniform over a multiset of pure strategies of size ds. We explain a strategy $(\mathcal{K}, \mathcal{T})$ for Player 1 that bounds the expected payoff. Player 1 chooses K by sampling d independent instances of the restriction of $\mathcal{Y}$ to S_0, and chooses T by sampling d independent instances of the restriction of $\mathcal{Y}$ to S_1. The expected payoff is

$$\begin{aligned}
E := &\Pr_{y \sim \mathcal{Y}}(y \in S_0)\ \mathbb{E}_{\pi,K,T}\left[\left\|R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_K^{\otimes (m_0+1)}, \mathcal{U}_T^{\otimes m_1}\right)\right)\right\| \middle| y \in S_0\right] \\
&+ \Pr_{y \sim \mathcal{Y}}(y \in S_1)\ \mathbb{E}_{\pi,K,T}\left[\left\|R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, y, \mathcal{U}_T^{\otimes m_1}\right)\right) - R\left(\pi\left(\mathcal{U}_K^{\otimes m_0}, \mathcal{U}_T^{\otimes (m_1+1)}\right)\right)\right\| \middle| y \in S_1\right] \ .
\end{aligned}$$

Let $x_1^{(0)}, x_2^{(0)}, \ldots, x_{d+1}^{(0)}$ and $x_1^{(1)}, x_2^{(1)}, \ldots, x_{d+1}^{(1)}$ be $d+1$ independent samples from $\mathcal{Y}|_{S_0}$ and $\mathcal{Y}|_{S_1}$, respectively. Sample $i^* \xleftarrow{\$} [d+1]$ and for $j \in \{0,1\}$, let $y_j^* := x_{i^*}^{(j)}$. Let $\widehat{\mathcal{Y}}_0$ and $\widehat{\mathcal{Y}}_1$ be the uniform distributions over the multisets $\{x_i^{(0)}\}_{i \in [d+1] \setminus \{i^*\}}$ and $\{x_i^{(1)}\}_{i \in [d+1] \setminus \{i^*\}}$, respectively. For $j \in \{0,1\}$, we have that $(y_j^*, \widehat{\mathcal{Y}}_0, \widehat{\mathcal{Y}}_1) \sim (\mathcal{Y}|_{S_j}, \mathcal{K}, \mathcal{T})$. Then, by Lemma A.5, we have

$$\mathbb{E}_\pi\left[\left\|R\left(\pi\left(\widehat{\mathcal{Y}}_0^{\otimes m_0}, y, \widehat{\mathcal{Y}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{Y}}_0^{\otimes (m_0+1)}, \widehat{\mathcal{Y}}_1^{\otimes m_1}\right)\right)\right\| \ \middle|\ y \in S_0\right] \ \leq \ \delta + \frac{2m_0+1}{d+1} \ ,$$

and

$$\mathbb{E}_\pi \left[\left\| R\left(\pi\left(\widehat{\mathcal{Y}}_0^{\otimes m_0}, y, \widehat{\mathcal{Y}}_1^{\otimes m_1}\right)\right) - R\left(\pi\left(\widehat{\mathcal{Y}}_0^{\otimes m_0}, \widehat{\mathcal{Y}}_1^{\otimes (m_1+1)}\right)\right) \right\| \;\middle|\; y \in S_1 \right] \;\leq\; \delta + \frac{2m_1+1}{d+1} \;.$$

Therefore, we obtain $E \leq \delta + 2(m+1)/(d+1)$.

Above, we showed that for every ds-uniform strategy for Player 2, there exists a strategy for Player 1 that bounds the expected payoff by $\delta + 2(m+1)/(d+1)$. Let $\mathbf{M} := [\mathbf{M}_{ij}]_{i,j}$ be the matrix such that $\mathbf{M}_{ij}$ corresponds to the payoff when Player 1 outputs i and Player 2 outputs j. By Lemma A.2, we have

$$\delta + 2(m+1)/(d+1) \geq \max_{\mathcal{Q} \in \mathfrak{Q}_{ds}} \min_i \mathbb{E}_{j \sim \mathcal{Q}}[\mathbf{M}_{ij}] \geq \omega(\mathbf{M}) - \varepsilon(\mathbf{M}_{\max} - \mathbf{M}_{\min}) \geq \omega(\mathbf{M}) - \varepsilon \;,$$

where $\mathfrak{Q}_{ds}$ is the set of all ds-uniform strategies for Player 2. It follows that $\omega(\mathbf{M}) \leq \delta + 2(m+1)/(d+1) + \varepsilon$.

Now we use Lemma A.2 in other way around. In fact, the number of possible choices for Player 1 is $|S_0 \cup S_1| \leq 2^n$. Therefore, Lemma A.2 asserts that there exists a s-uniform strategy for Player 2 such that for any possibly mixed strategy for Player 1, the expected payoff is at most ε-far from the value of the game $\omega(\mathbf{M})$. In other words, for this particular strategy of Player 1, the expected payoff is always at most

$$\omega(\mathbf{M}) + \varepsilon \leq \delta + 2(m+1)/(d+1) + 2\varepsilon \,.$$

Recall that a s-uniform strategy is, by definition, a uniformly sampled element from a size-s multiset of choices of the player. Note that Player 1 chooses a pair (K, T). Therefore, this strategy is essentially a uniform distribution over some multiset $\{(K_1, T_1), \ldots, (K_s, T_s)\}$, which concludes the proof.

□

Table of Contents

Amplifying Randomized Encodings & Applications 1
Pouria Fallahpour, Alex B. Grilo, Garazi Muguruza , and Mahshid Riahinia
1 Introduction 1
1.1 Our Contribution 2
Application 1: Zero-knowledge Amplification of NISZK. 4
Application 2: OWFs and the Worst-Case Hardness of SZK. 4
Application 3: OWFs from Weak and Imperfect iO. 5
Application 4: One-Way State Generators. 6
1.2 Technical Overview 7
1.3 Other Related Works 11
1.4 Open Questions 11
2 Preliminaries 11
3 Complexity Classes 16
4 Sparsely Lossy Reductions Are Randomized Encodings 18
4.1 Better Analysis for Small λ 21
5 Randomized Encodings Are Sparsely Lossy 22
6 Randomized Encodings and One-Wayness 25
7 Obfuscation 28
A Proof of Extended Disguising Lemma 33